Jet Propulsion Laboratory
California Institute of Technology

Agile Science

A New Paradigm for Space Missions

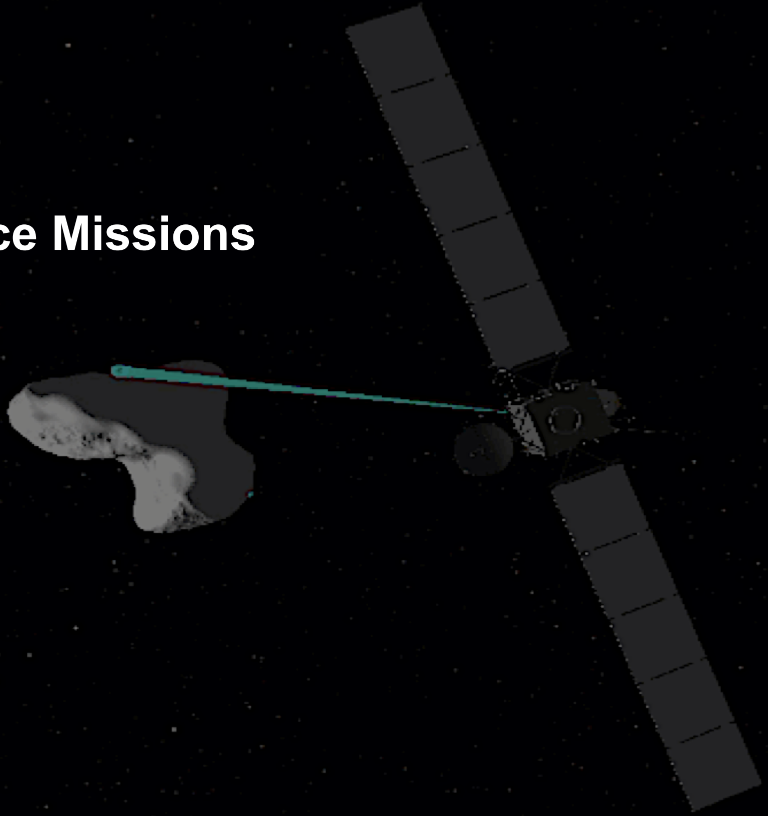

Jack Lightholder
David Brown
Steve Chien

Brian Bue (398J)
Julie Castillo-Rogez (3227)
William Huffman (398J)
Gregg Rabideau (397I)
Umma Rebbapragada (398J)
David R. Thompson (382B)
Martina Troesch (397I)
Kiri Wagstaff (398J)

This research was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The Agile Science Paradigm

Analyze data acquired onboard spacecraft and respond based on analysis

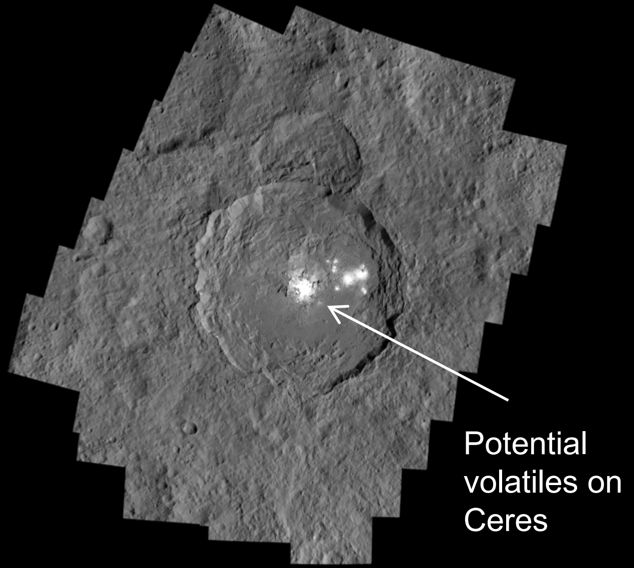

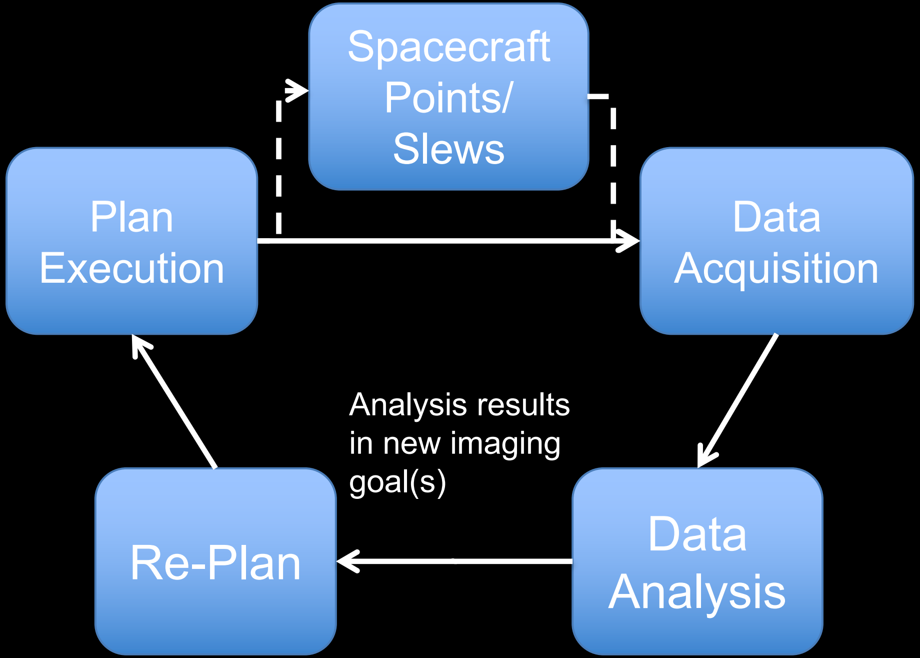

Image Credit: NASA/JPL/Dawn Mission

Agile Science in Orbit

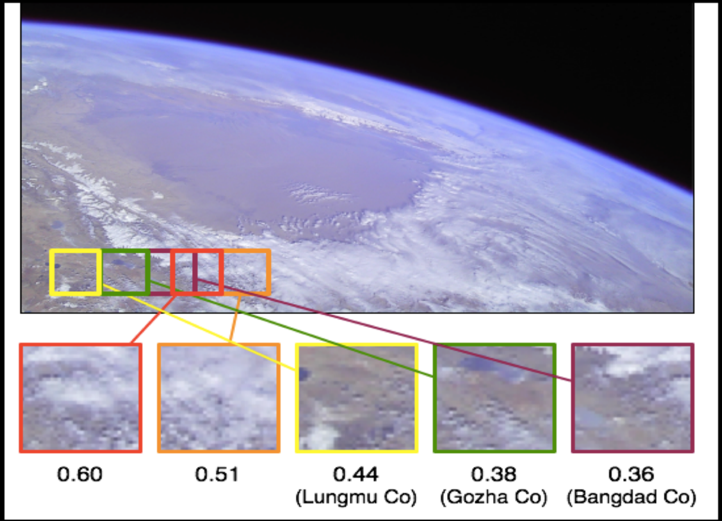

Visual Salience: Identified areas of the image that differ from surrounding areas.

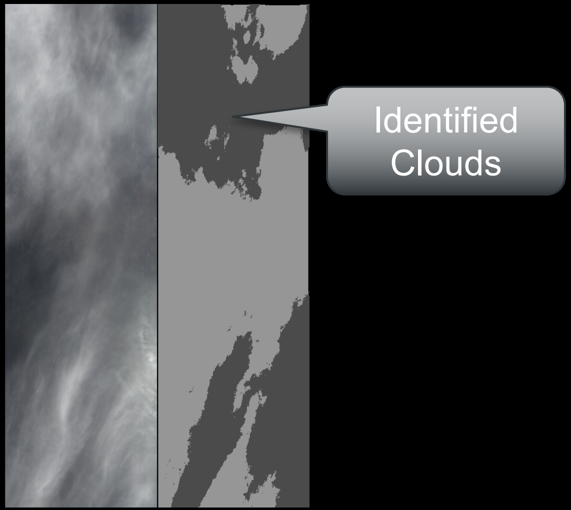

Preliminary Cloud Classification results from EO-1

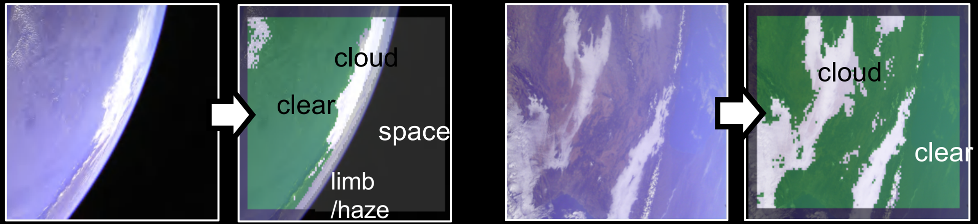

TextureCam: Pixel classification for cloud screening, downlink prioritization

[Chien et al. 2016 JAIS; Thompson et al., i-SAIRAS 2012; Wagstaff et al., *GRL* 2013; Bekker et al., *Astrobiology* 2014; Altinok et al. JFR 2015, Thompson et al. TGARS 2011]

Agile Science Landed

Autonomous Exploration for Gathering Increased Science (AEGIS)

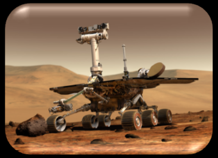

Mars Exploration Rover (2009)

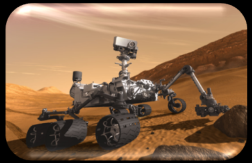

Mars Science Laboratory (2012)

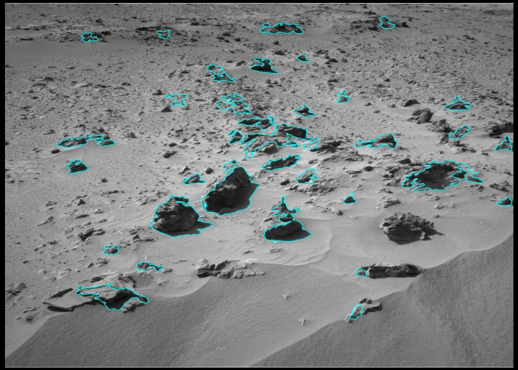

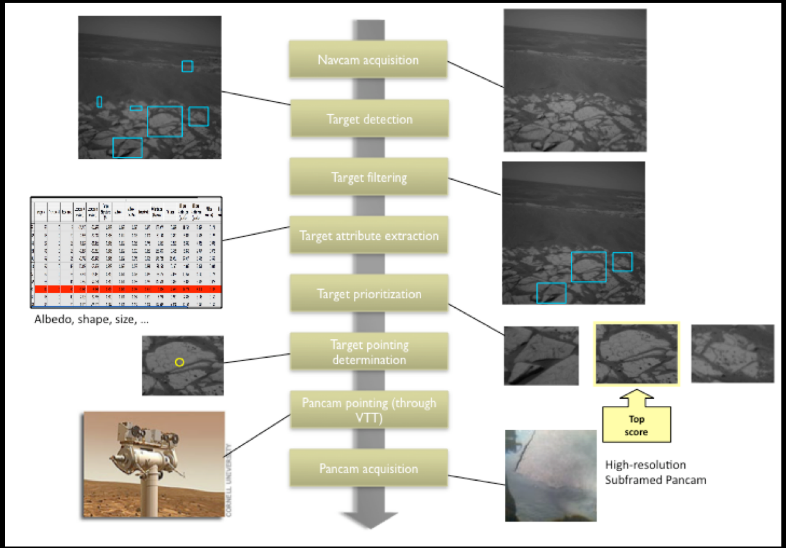

- Provides **intelligent targeting and data acquisition** by:
  - analyzing images of the rover scene
  - identifying high-priority science targets (e.g., rocks)
  - taking follow-up imaging of these targets with no ground communication required

[Estlin et al. 2011]

A Deep Space Example

Near Earth Asteroid Scout

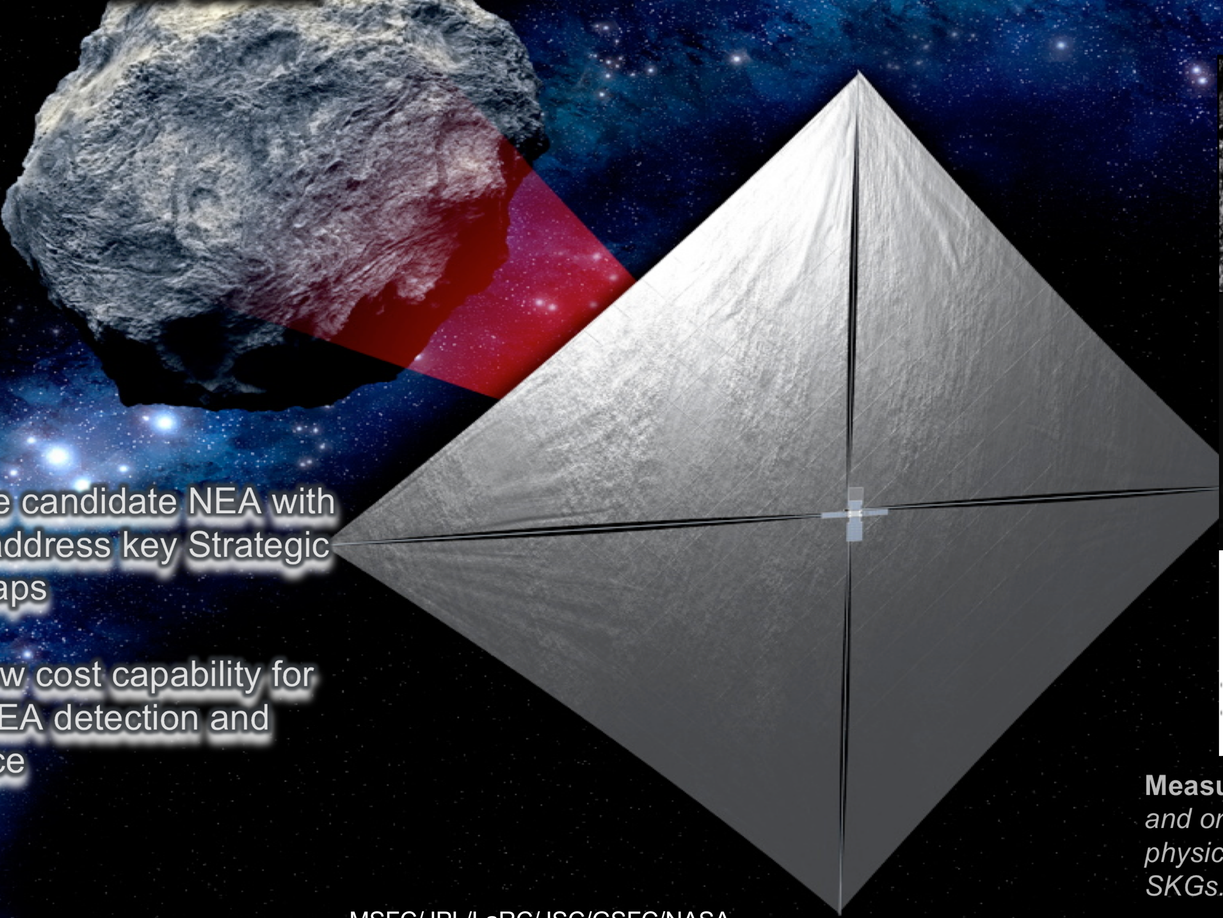

GOALS

Characterize one candidate NEA with an imager to address key Strategic Knowledge Gaps

Demonstrates low cost capability for HEOMD for NEA detection and reconnaissance

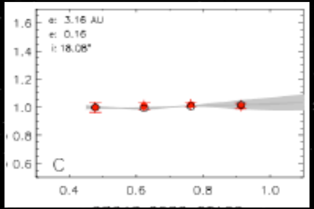

Measurements: *NEA volume, spin and orbital properties, address key physical and regolith mechanical SKGs.*

MSFC/JPL/LaRC/JSC/GSFC/NASA

G–66840

Imaging Challenges

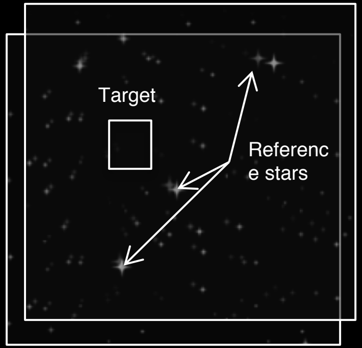

Target Detection and Approach
Ephemeris determination

Target Position Uncertainty

Spacecraft Pointing and Camera Limitations

Medium Field Imaging
Shape, spin, and local environment

Short Flyby Time (<30 minutes)

Uncertain Environment

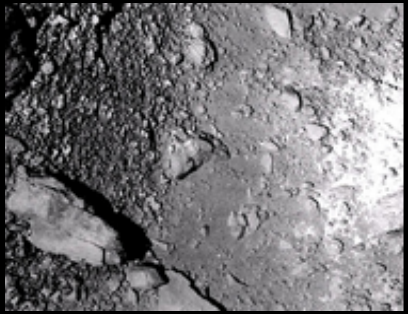

Close Proximity Imaging
Local scale morphology, terrain properties

Data Value Analysis and Sorting

Short Time at Closest Approach (<10 minutes)

Limited Downlink of 1 Kbps

Raw Data is Messy

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Raw Data is Messy

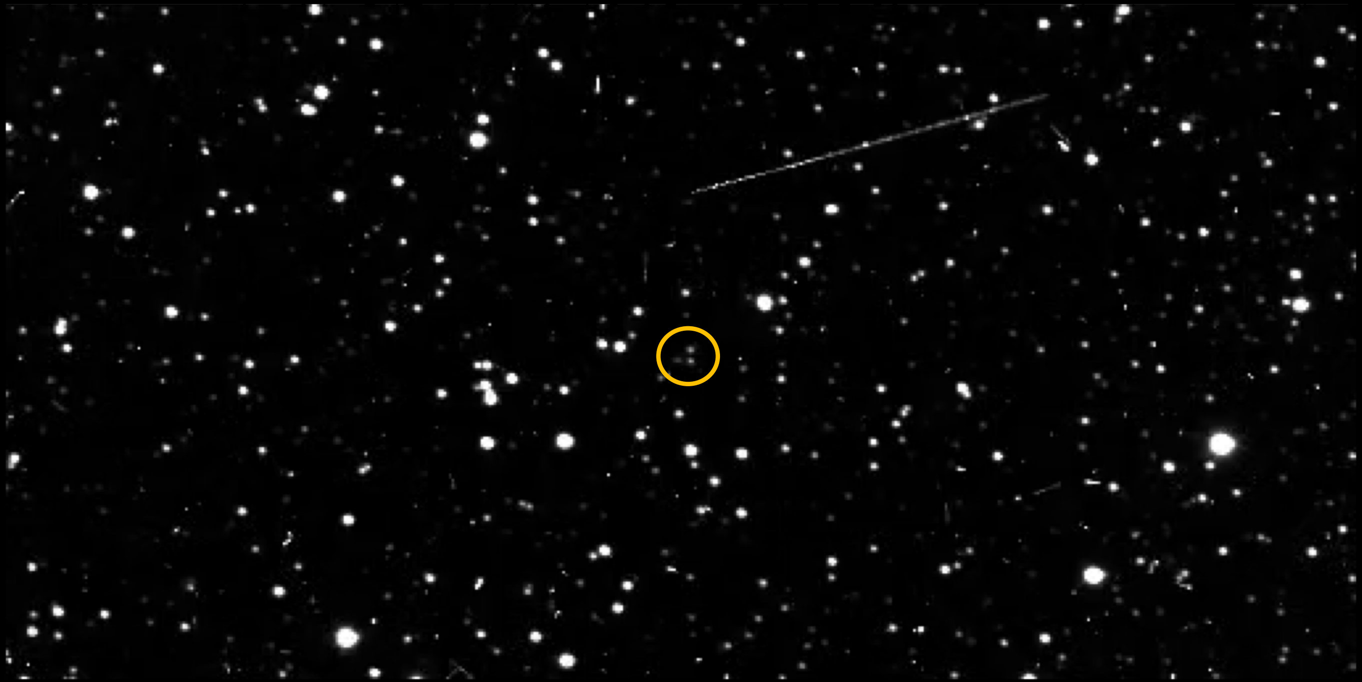

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Mission Operations Flexibility

Data Calibration

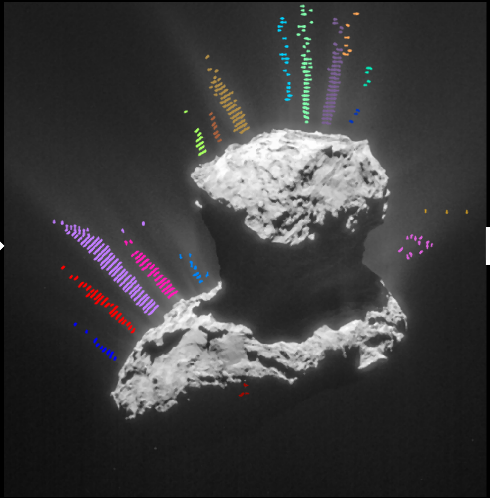

Feature Detection

Operational Decisions

Mission Operations Flexibility

Calibration

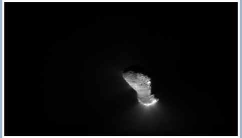

- Dark Current Subtraction
- Flat field / Radiometric calibration

Onboard Feature Detection

- Reference Star Detection
- Target Detection
- Cosmic Ray Removal

Follow-Up Actions

- Prioritize Data for Downlink
- Downlink “Quick-Look” Data Products

Cleaning Up the Noise

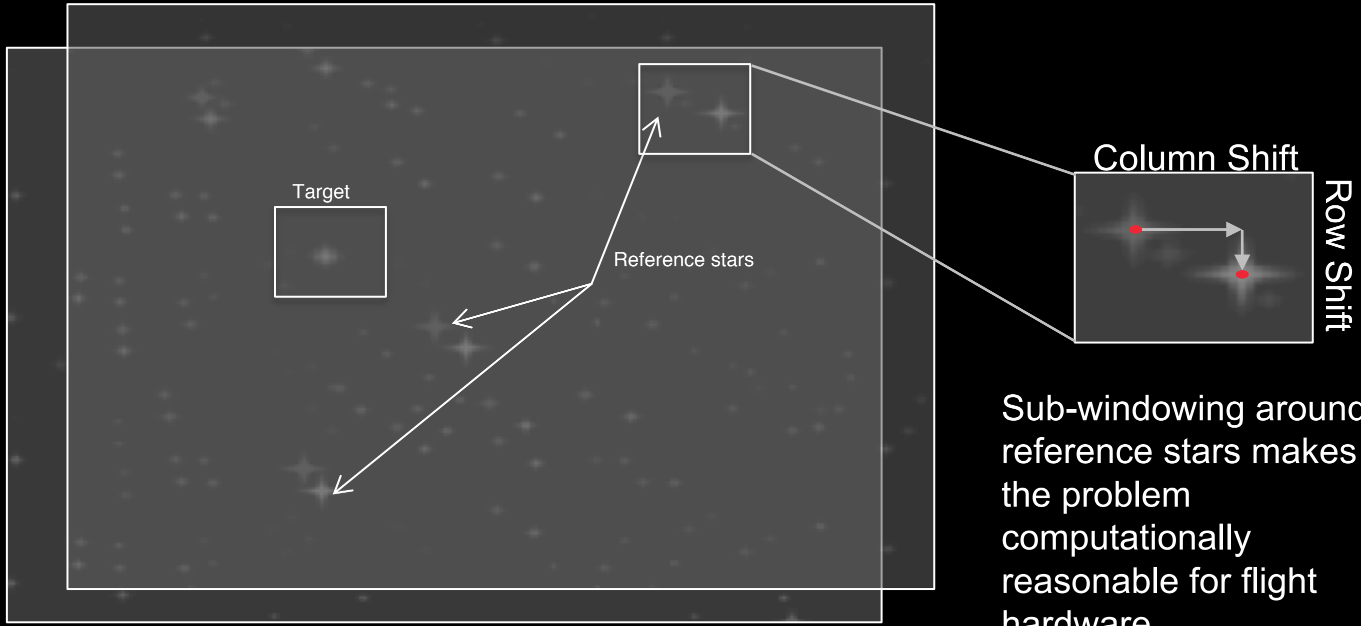

Sub-windowing around reference stars makes the problem computationally reasonable for flight hardware.

Onboard co-registration of images improves SNR and reduces downlink requirements

[Thompson et al. 2015]

Computation is additionally constrained by onboard memory limitations.

< 100 MB RAM

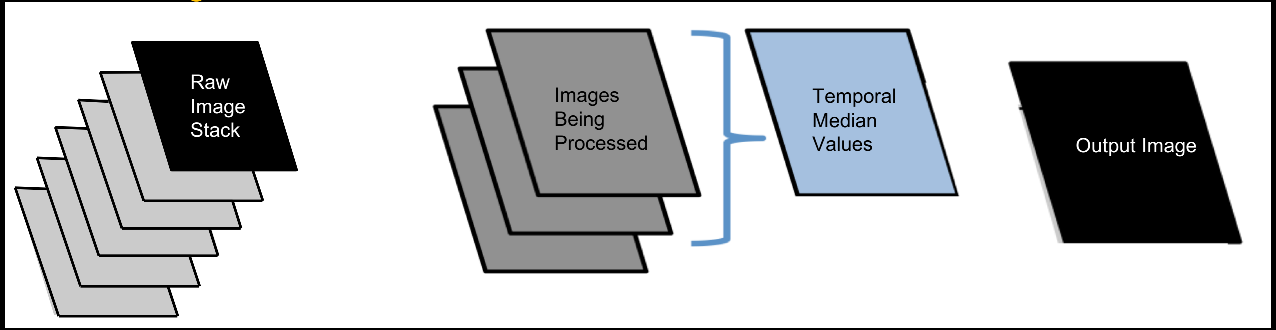

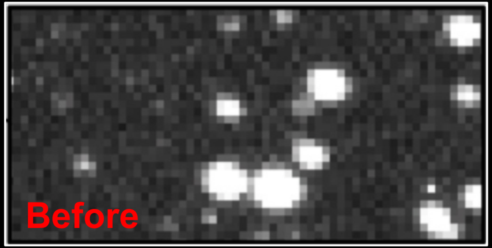

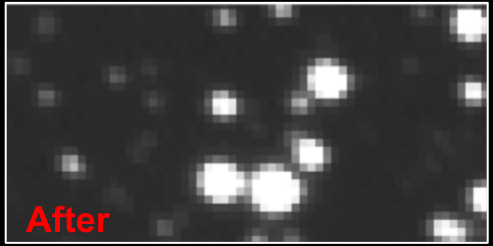

Stepwise processing keeps the necessary memory small.

[Thompson et al. 2015]

Using data from multiple images improves data quality.

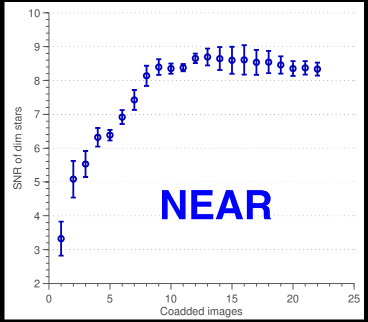

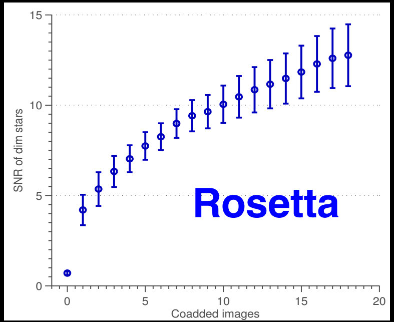

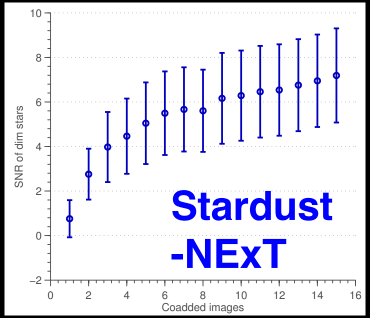

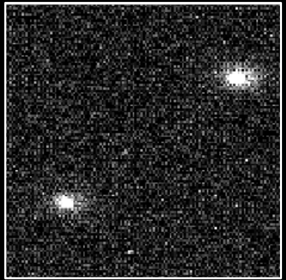

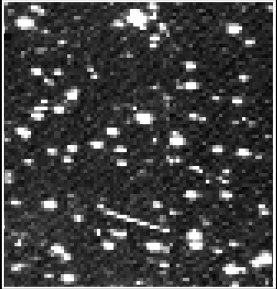

Save time and bandwidth while improving situational awareness.

[Thompson et al. 2015]

Processed Data

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Identify Targets with Onboard Image Subtraction

Determine the shift between two images, subtract with (x,y) offset.

This type of information has many mission applications.

Current trajectory verification and refinement

Automated target tracking

Target of opportunity detection

Target survey and classification

Does Your Target Look "As Expected"?

New Horizons Long Range Reconnaissance Imager Detection of Pluto/Charon

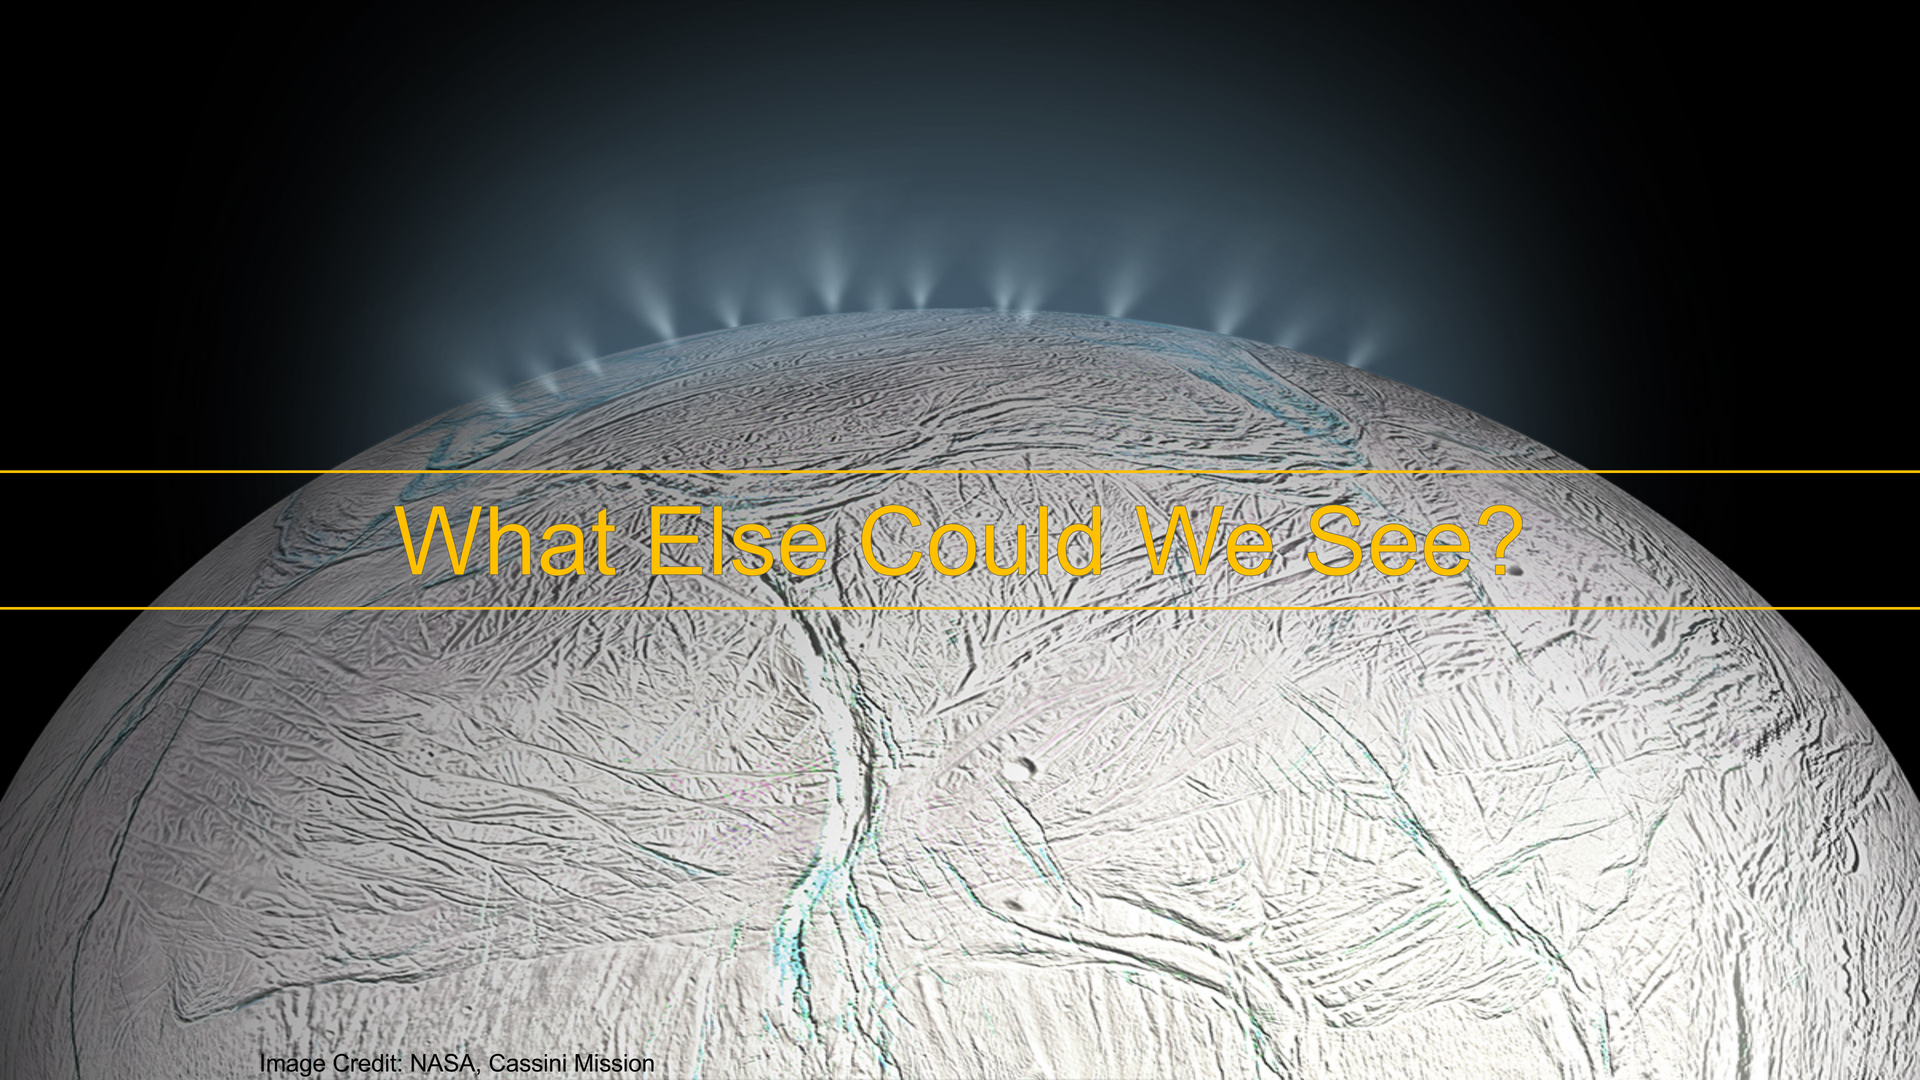

Image Credit: NASA, Cassini Mission

Plumes are Scientifically Exciting

Enceladus

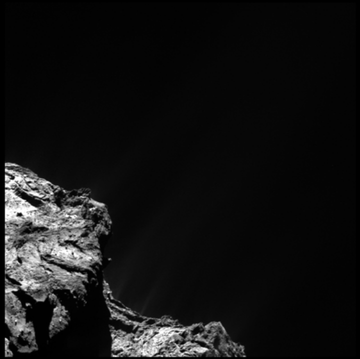

Comet 67P

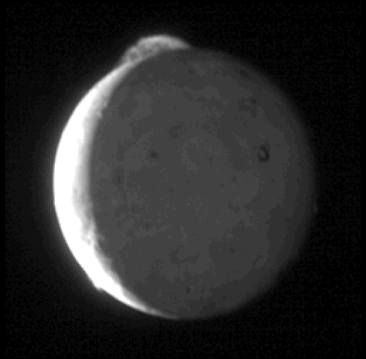

Io

Plumes gives scientists insights into the volatiles located throughout the solar system.

Unfortunately, they're not scheduled. We have to react fast.

Image Credit: NASA, Cassini Mission (Left), ESA, Rosetta Mission (Center), NASA, New Horizons Mission(Right)

Plume Detection

Hartley 2 (*EPOXI*)

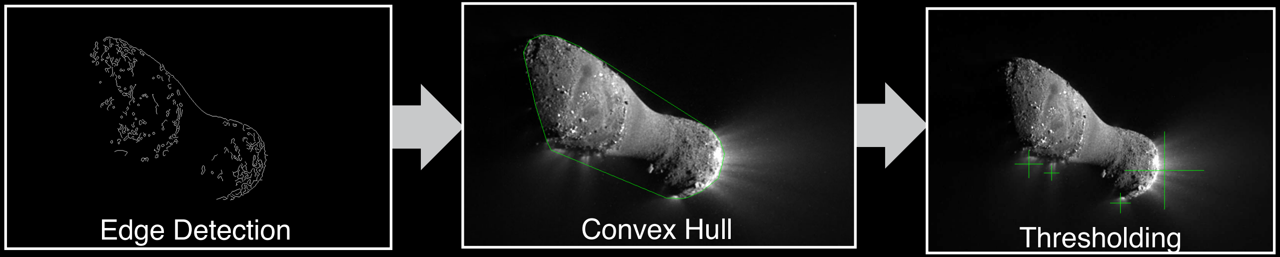

- Detects bright material beyond the limb
- Enables monitoring campaigns, target-relative data acquisition
- Detects most plumes with zero false positives

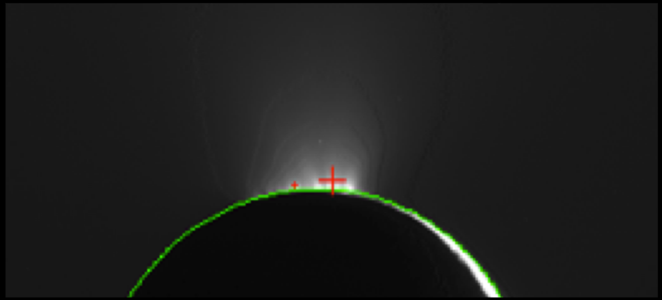

Enceladus (*Cassini*)

[Thompson et al., PSS 2012, Wagstaff et al, ApJ 2014]

Comet Tracking

Hartley 2 flyby
Original Sequence

Agile Science Planning

Target and plume detection using MRI-VIS on EPOXI. Tracking Comet Hartley 2.

Jet Propulsion Laboratory
California Institute of Technology

Churyumov Gerasimenko (C-G) Plume Detection

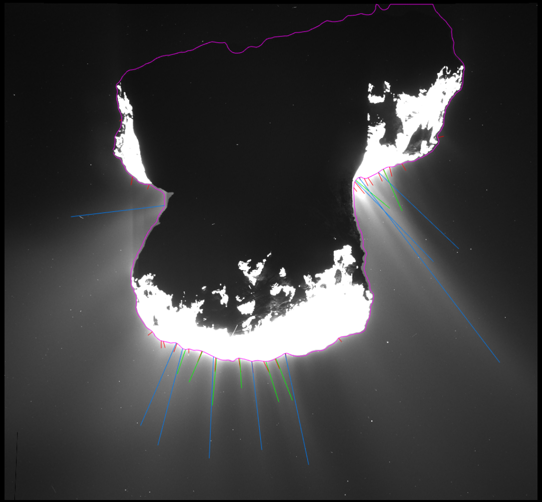

David Brown[1]
Steve Chien[1]
Sierks Holger[2]
William Huffman[1]
David Thompson[1]
Jean-Baptiste Vincent[2]

[1] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA
[2] Max-Planck-Institut für Sonnensystemforschung, Göttingen, Germany

Portions of this research were carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Motivation

- C-G plumes are an ideal agile science candidate

Motivation

- C-G plumes are an ideal agile science candidate
- High value events
  - Insight into the makeup of early solar system body

Motivation

- C-G plumes are an ideal agile science candidate
- High value events
  - Insight into the makeup of early solar system body
- Transient events
  - Not long enough to do ground-based instrument planning

Motivation

- C-G plumes are an ideal agile science candidate
- High value events
  - Insight into the makeup of early solar system body
- Transient events
  - Not long enough to do ground-based instrument planning
- Unpredictable events

Motivation

- C-G plumes are an ideal agile science candidate
- High value events
  - Insight into the makeup of early solar system body
- Transient events
  - Not long enough to do ground-based instrument planning
- Unpredictable events
- Small field-of-view instruments (e.g. MIRO, ALICE) mandate high-precision pointing
  - Broad sweeps of comet body fail to meet some science goals

Broad Sweeps vs Targeted Sweeps

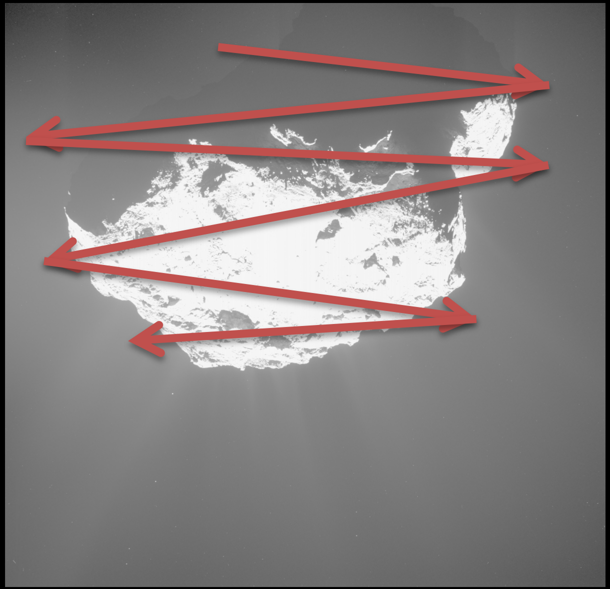

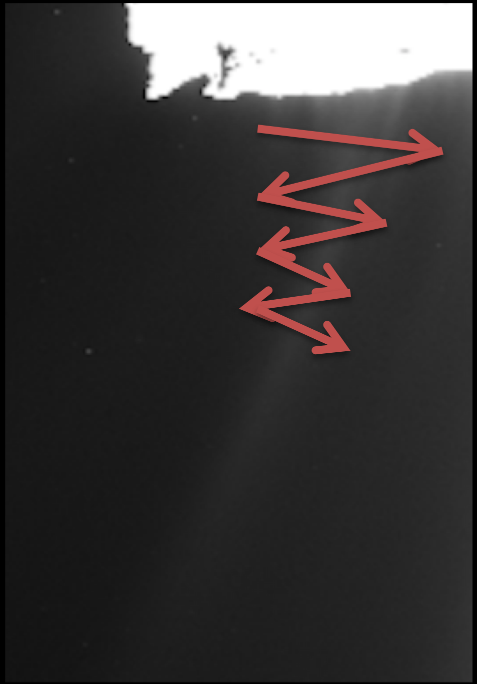

Motivation

- C-G plumes are an ideal agile science candidate
- High value events
  - Insight into the makeup of early solar system body
- Transient events
  - Not long enough to do ground-based instrument planning
- Unpredictable events
- Small field-of-view instruments (e.g. MIRO, ALICE) mandate high-precision pointing
  - Broad sweeps of comet body fail to meet some science goals
- Translatable to other mission applications

Challenges

- High contrast environment
- Noisy images
- Uncertain pointing
- Variability in plume profiles

Challenges – High Contrast Environment

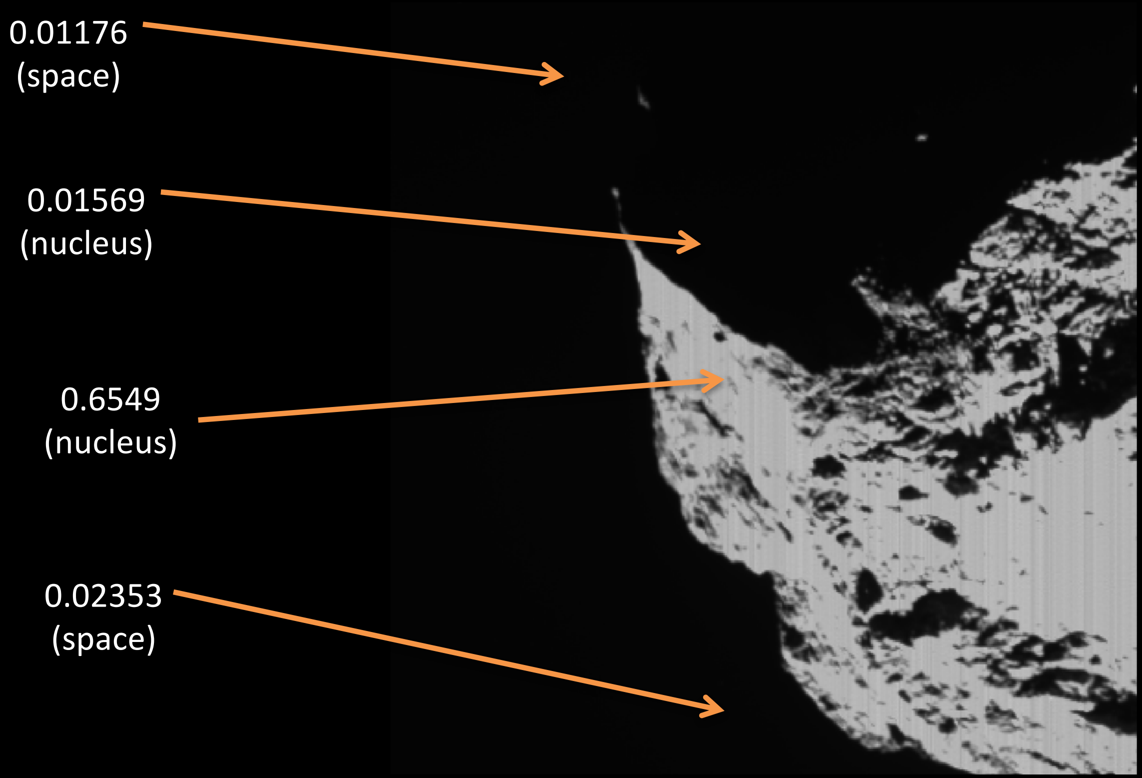

Challenges

- High contrast environment
- Variability in plume profiles
- Noisy images
- Uncertain pointing

Challenges – Variability in Plume Profiles

Long, narrow, curved

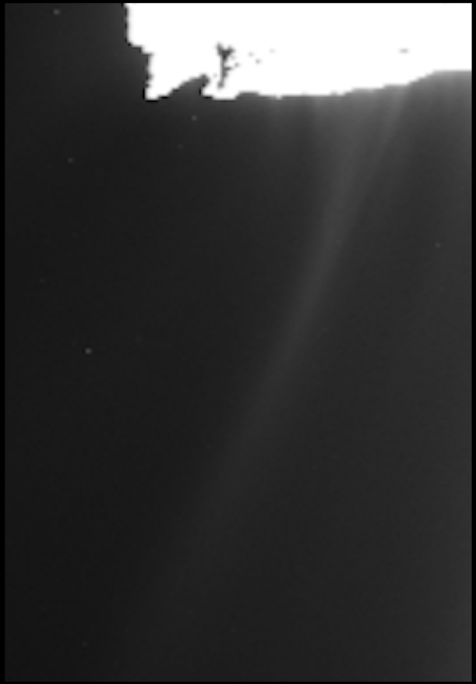

Short, diffuse, straight

Challenges

- High contrast environment
- Variability in plume profiles
- Noisy images
- Uncertain pointing

How to find a plume

- Using visual information, we are limited to hunting for plumes that extend beyond the nucleus

- Plume detection pipeline:
  1. Outline the nucleus body

How to find a plume

- Using visual information, we are limited to hunting for plumes that extend beyond the nucleus

- Plume detection pipeline:
  1. Outline the nucleus body
  2. Detect candidate plume features in the image

How to find a plume

- Using visual information, we are limited to hunting for plumes that extend beyond the nucleus

- Plume detection pipeline:
  1. Outline the nucleus body
  2. Detect candidate plume features in the image
  3. Convert into plume vectors

How to find a plume

- Using visual information, we are limited to hunting for plumes that extend beyond the nucleus

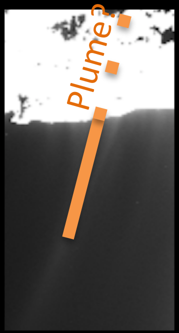

- Plume detection pipeline:
  1. Outline the nucleus body
  2. Detect candidate plume features in the image
  3. Convert into plume vectors
  4. Correlate across time

Plume Detection Pipeline

1. Outline the nucleus body
2. Detect candidate plume features in the image
3. Convert into plume vectors
4. Correlate across time

Finding the Nucleus - Nucleus Projection

- Start with projection from nucleus shape model

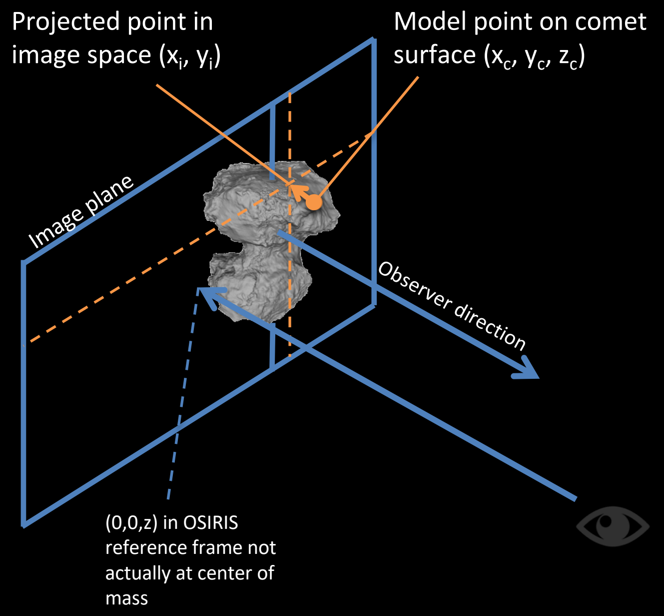

Finding the Nucleus - Nucleus Projection Errors

- Errors in projection accuracy
  - Navigation data
  - Projection errors due to flattening assumption (at close approach)
- Worst observed case: 70 pixels at 30km = ~200m

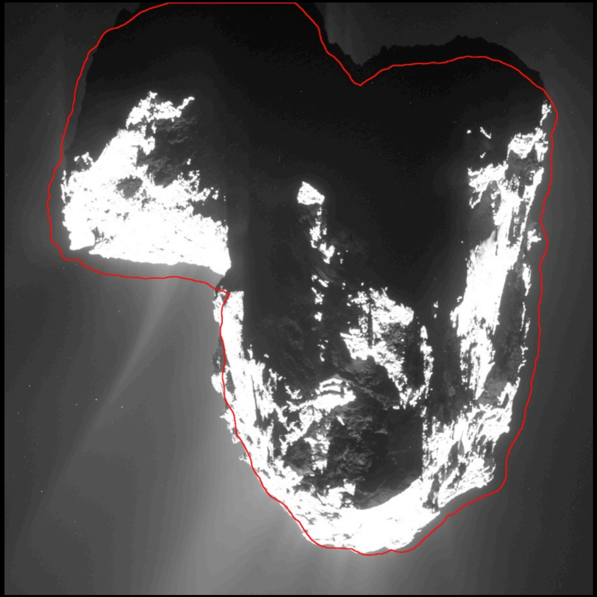

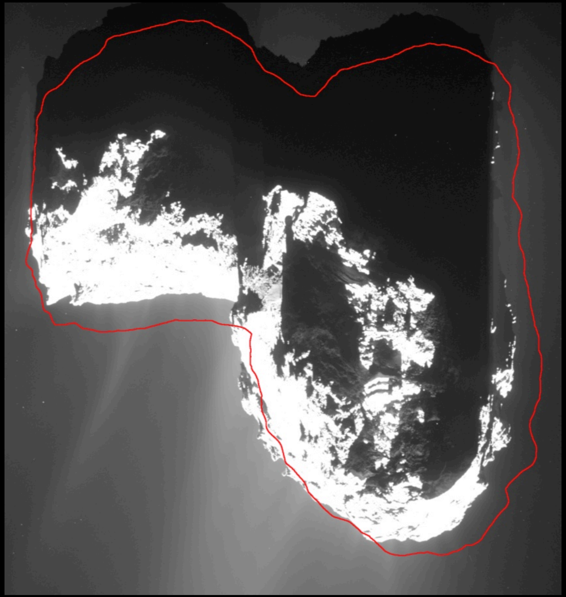

Plume Detection Pipeline

1. Outline the nucleus body
    a. Generate shape model projection
    b. Refine projection via image-based high fidelity tracing
2. Detect candidate plume features in the image
3. Convert into plume vectors

Plume Detection Pipeline

1. Outline the nucleus body
    a. Generate shape model projection
    b. Refine projection via image-based high fidelity tracing
        i. Prepare images for edge detection
        ii. Edge detection
        iii. Segment image into nucleus and non-nucleus
2. Detect candidate plume features in the image
3. Convert into plume vectors

Finding the Nucleus – Image Pre-Processing

- Log transform

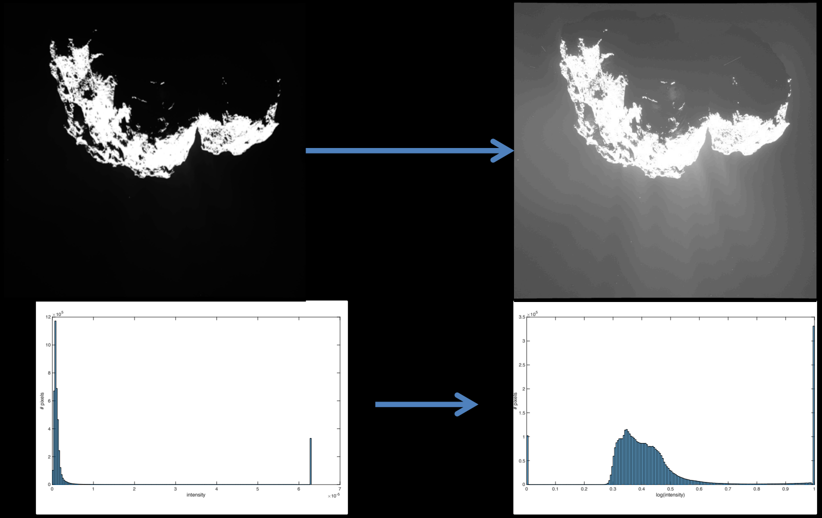

- Median filter
  - Stars and cosmic rays

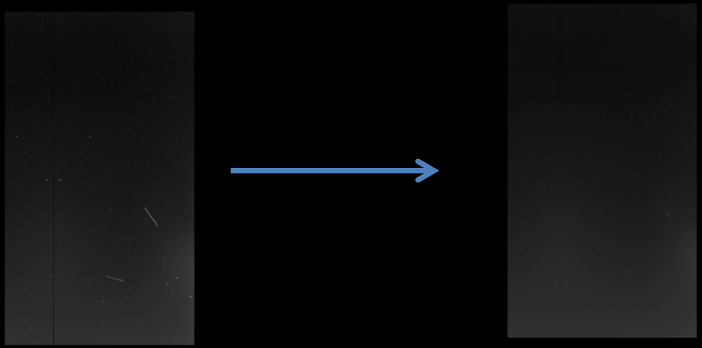

Finding the Nucleus – Edge Accentuation

1. Image Variance

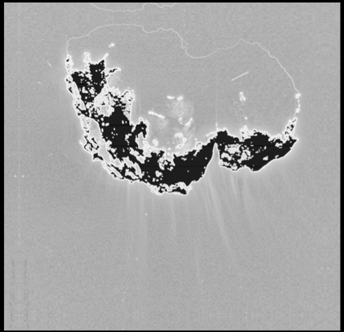

2. Relative Variance

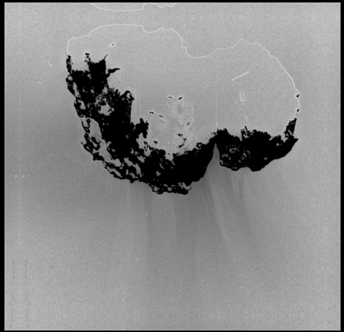

3. Gradient

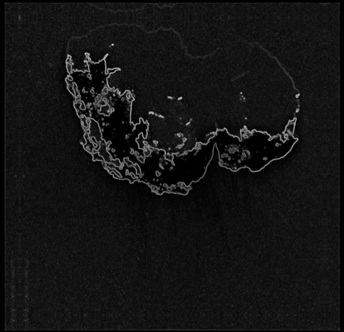

4. Combine

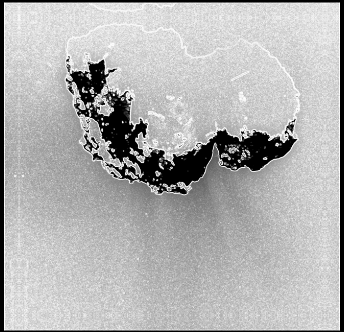

5. Filter

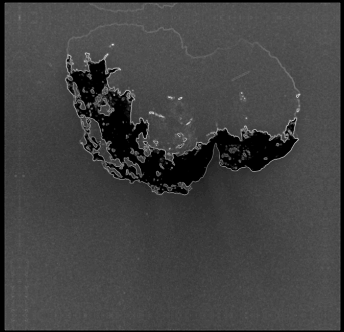

6. Gradient

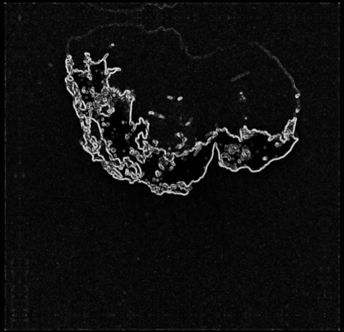

Plume Detection Pipeline

1. Outline the nucleus body
   a. Generate shape model projection
   b. Refine projection via image-based high fidelity tracing
      i. Prepare images for edge detection
      ii. Edge detection
      iii. Segment image into nucleus and non-nucleus
2. Detect candidate plume features in the image
3. Convert into plume vectors

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. Shrink contour inwards using active contours (Chan-Vese [1])
   - Energy function pulls contour inwards, and edges trigger resistance

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. Shrink contour inwards using active contours (Chan-Vese [1])
   - Energy function pulls contour inwards, and edges trigger resistance

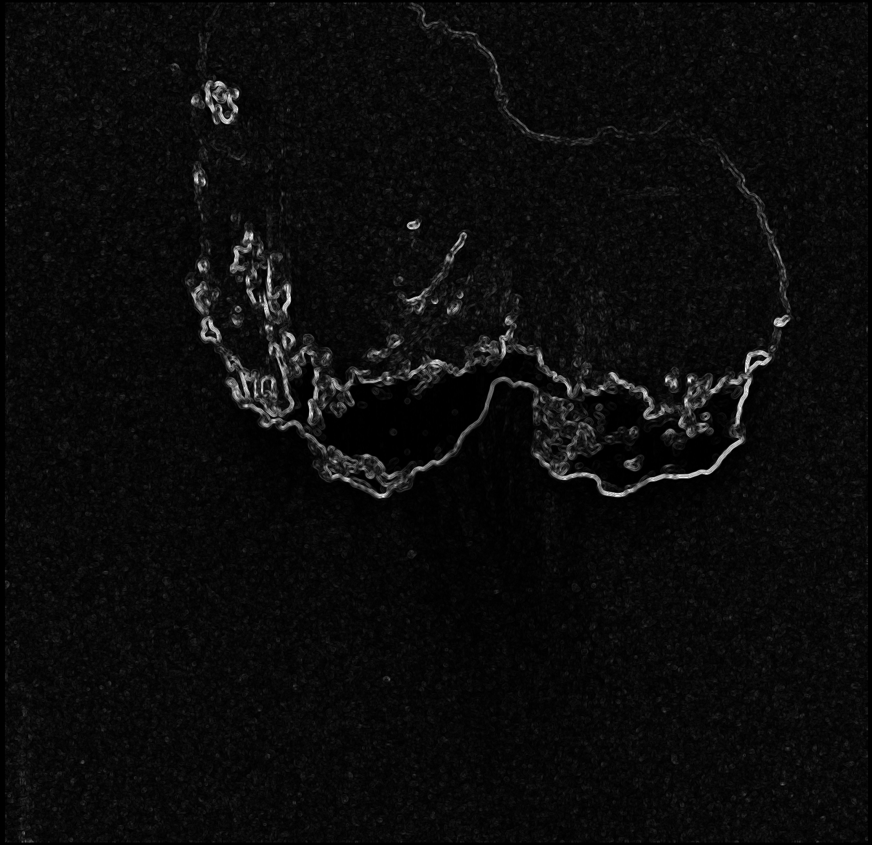

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. Shrink contour inwards using active contours (Chan-Vese [1])
   - Energy function pulls contour inwards, and edges trigger resistance

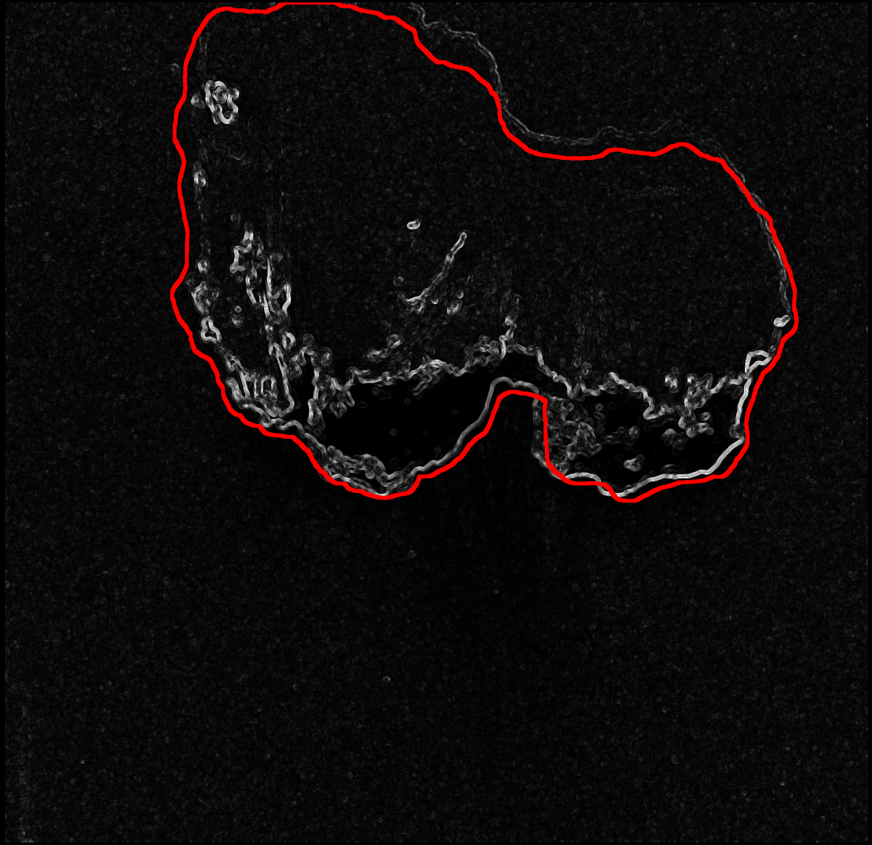

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. Shrink contour inwards using active contours (Chan-Vese [1])
   - Energy function pulls contour inwards, and edges trigger resistance

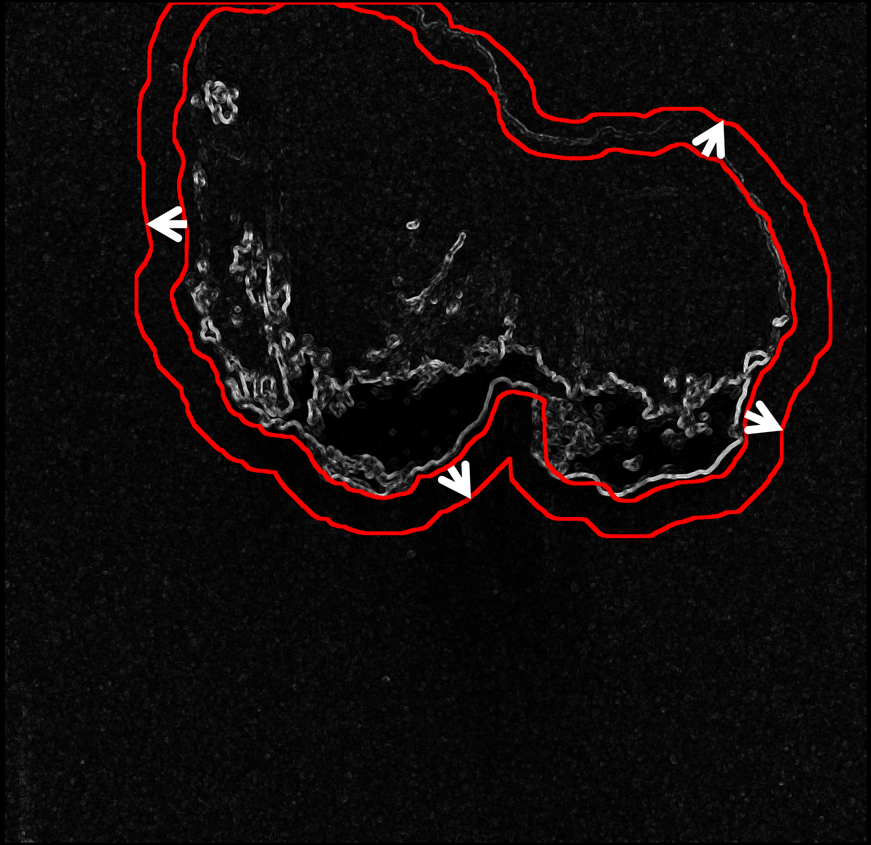

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. Shrink contour inwards using active contours (Chan-Vese [1])
   - Energy function pulls contour inwards, and edges trigger resistance

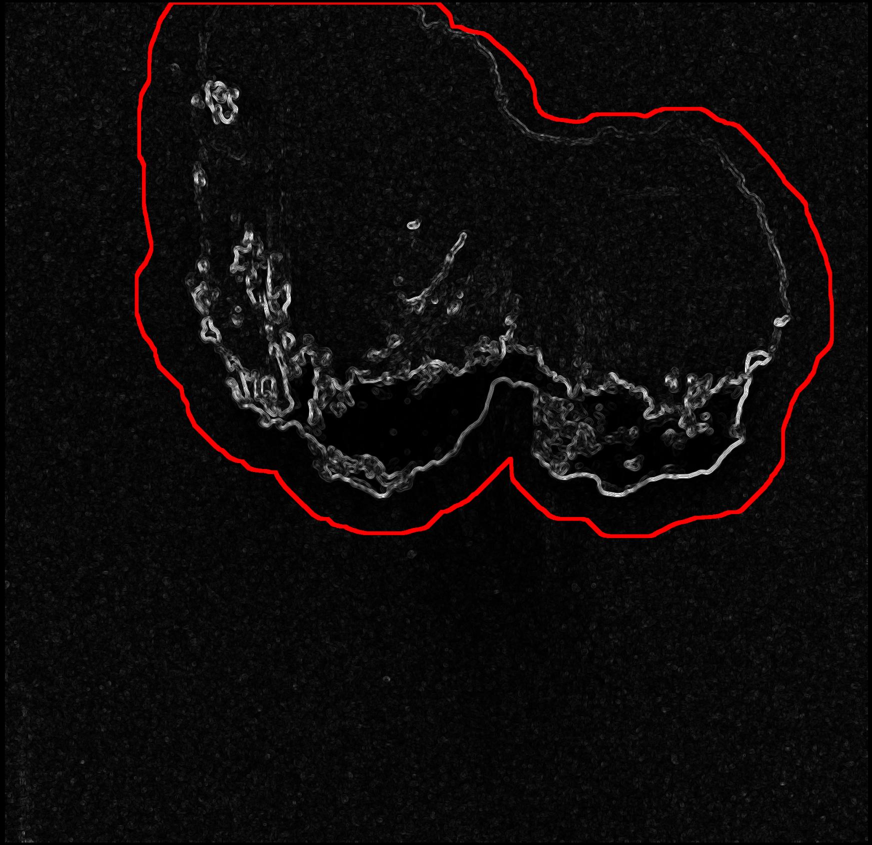

Finding the Nucleus – Image Segmentation

1. Take edge-accentuated image
2. Begin with shape model projection
3. Dilate based on expected error rate
4. **Shrink contour inwards using active contours (Chan-Vese [1])**
   - Energy function pulls contour inwards, and edges trigger resistance

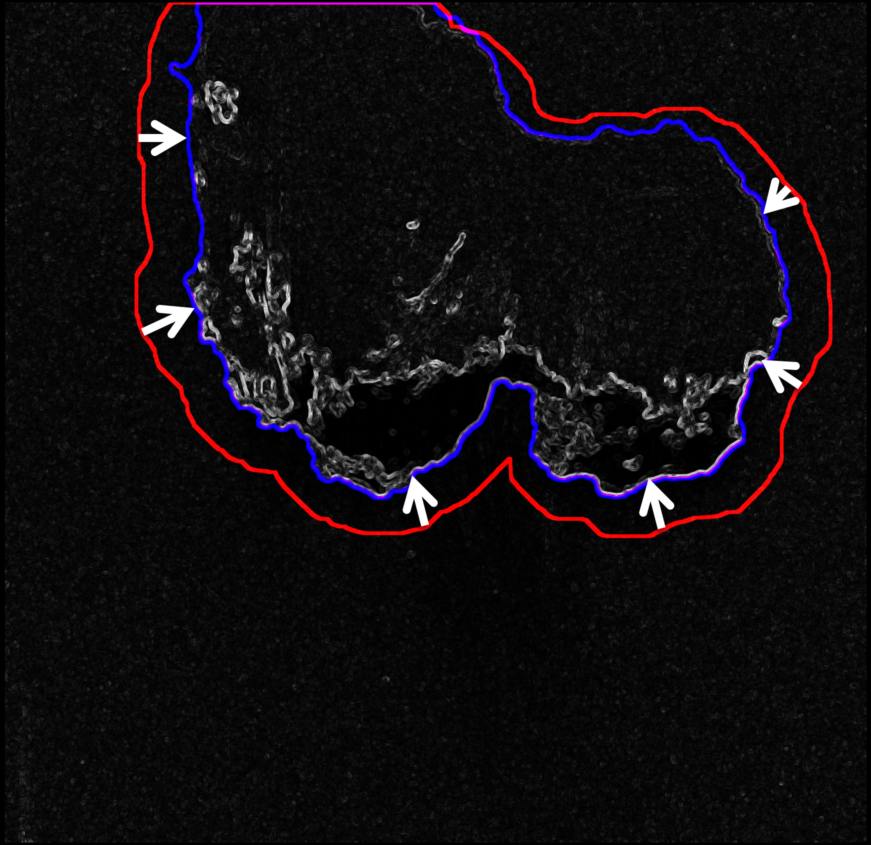

Finding the Nucleus – Image Segmentation

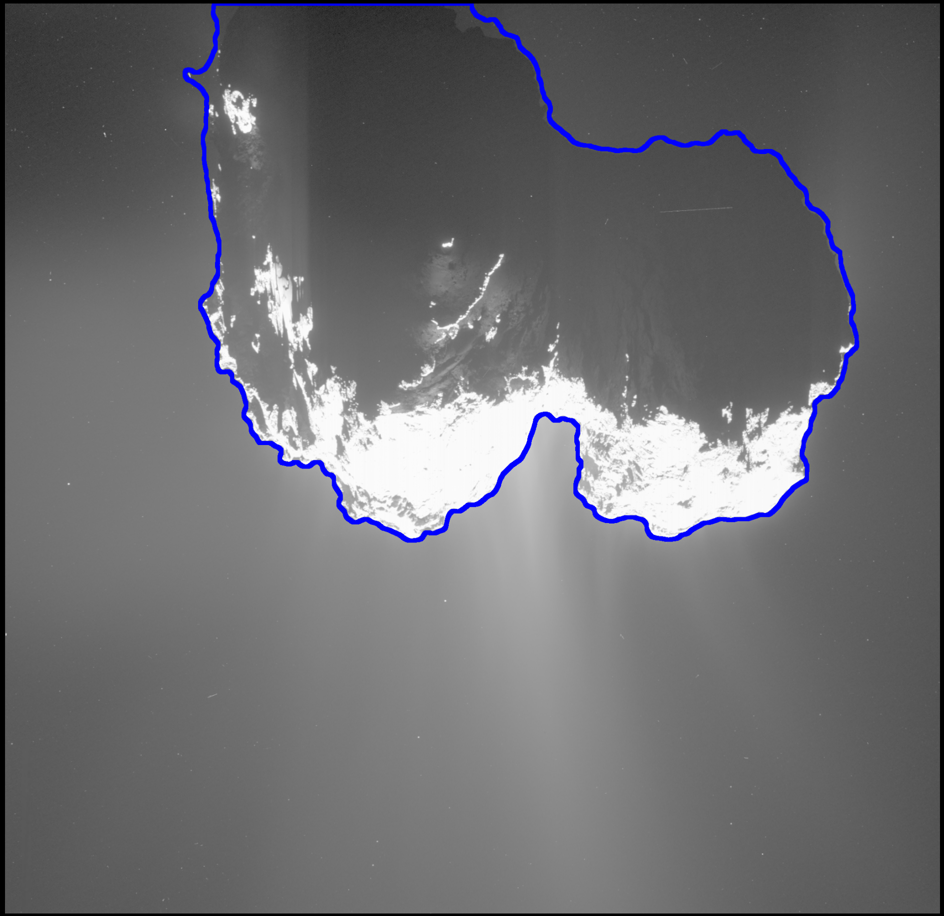

Plume Detection Pipeline

1. Outline the nucleus body
2. Detect candidate plume features in the image
3. Convert into plume vectors
4. Correlate across time

Plume Detection – Plume Candidate Areas

- Plumes are narrow streams of particles that are illuminated against the background
- Expect to see a region of higher intensity than its surroundings
- For each region, compare a narrow median filter with a wide median filter

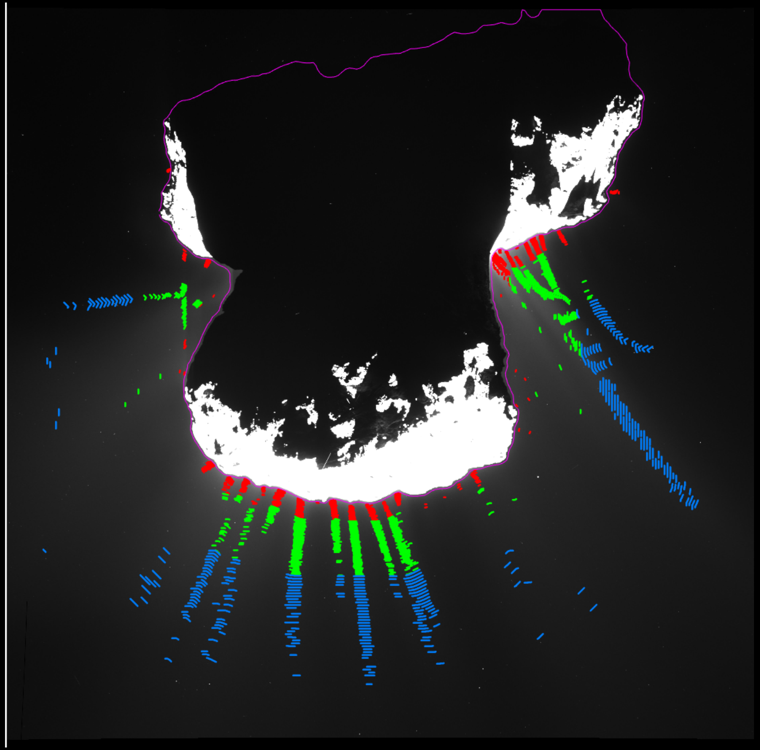

Plume Detection – Plume Candidate Areas

- Plumes are narrow streams of particles that are illuminated against the background
- Expect to see a region of higher intensity than its surroundings
- For each region, compare a narrow median filter with a wide median filter

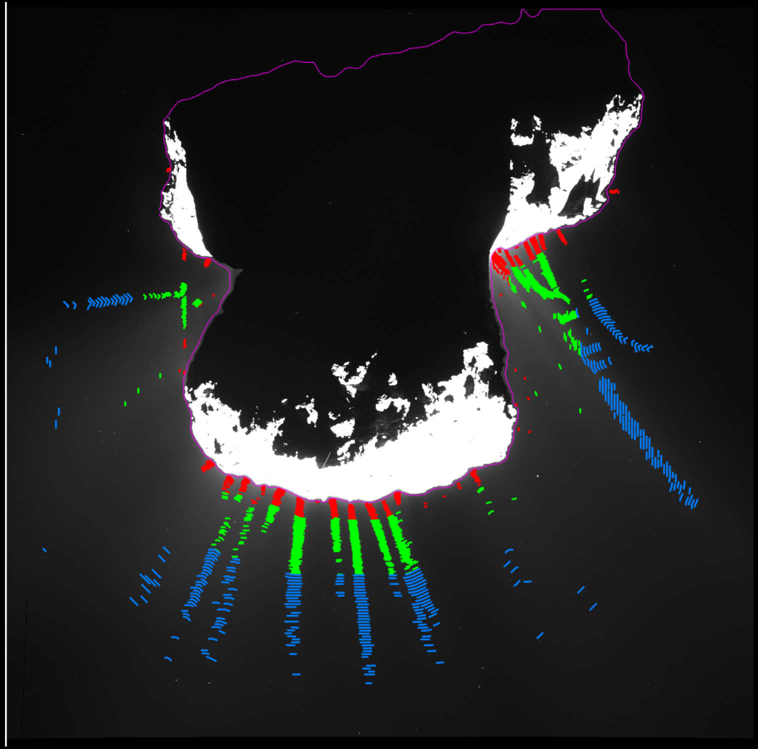

Plume Detection – Plume Candidate Areas

- Plumes are narrow streams of particles that are illuminated against the background
- Expect to see a region of higher intensity than its surroundings
- For each region, compare a narrow median filter with a wide median filter

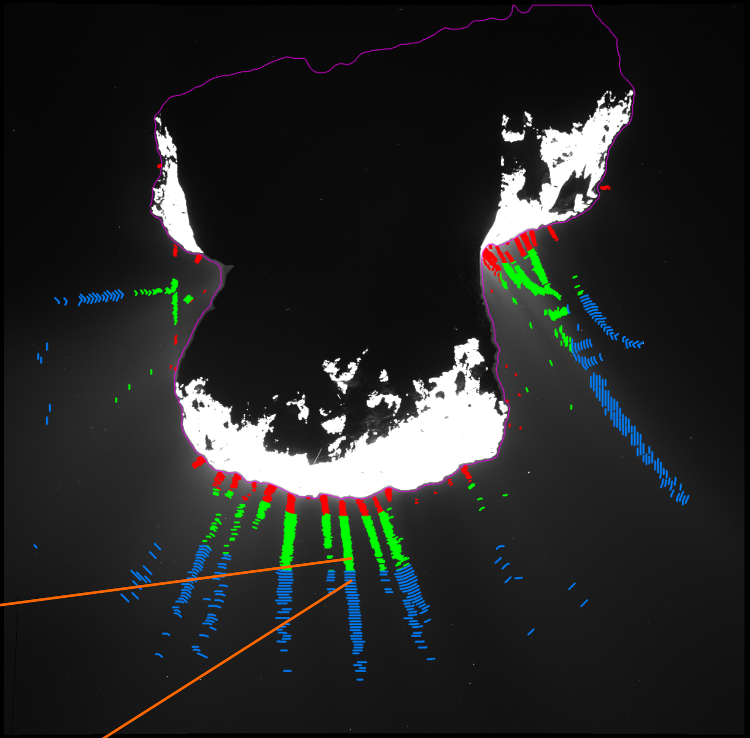

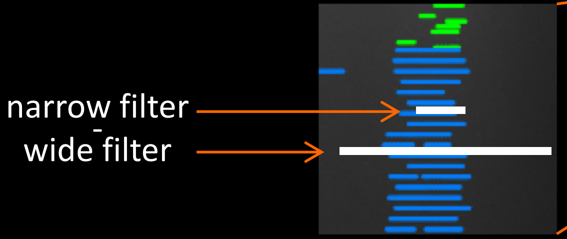

Plume Detection Pipeline

1. Outline the nucleus body
2. Detect candidate plume features in the image
3. Convert into plume vectors
4. Correlate across time

Plume Detection – Plume Vectorization

1. Group nearby pixels together
2. Run Random Sample Consensus (RANSAC) robust line-fitting algorithm on each group, weighted by intensity
3. Apply prior knowledge constraints:
   - Assume plumes are more normal to the nucleus contour
   - Extending plume should intersect nucleus contour

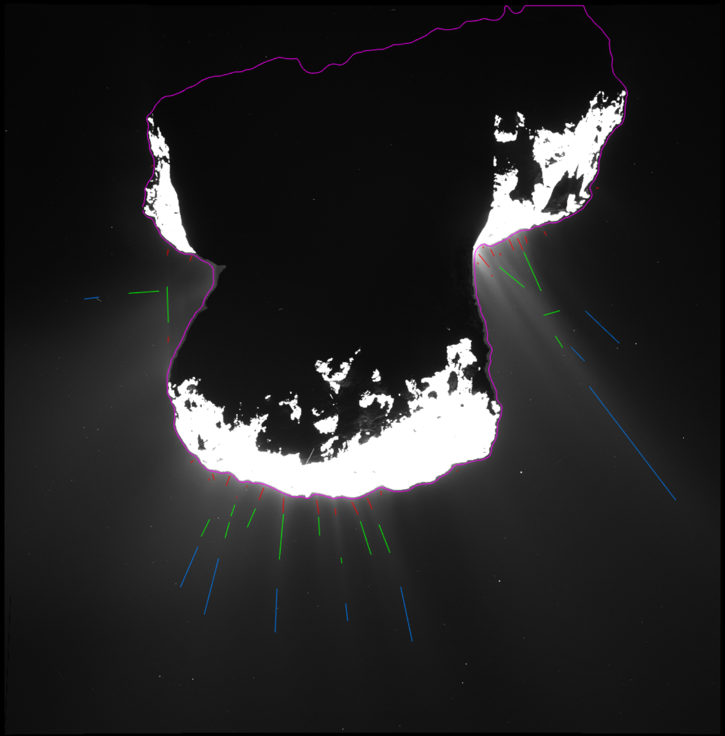

Plume Detection – Plume Vectorization

1. Group nearby pixels together
2. Run Random Sample Consensus (RANSAC) robust line-fitting algorithm on each group, weighted by intensity
3. Apply prior knowledge constraints:
   - Assume plumes are more normal to the nucleus contour
   - Extending plume should intersect nucleus contour

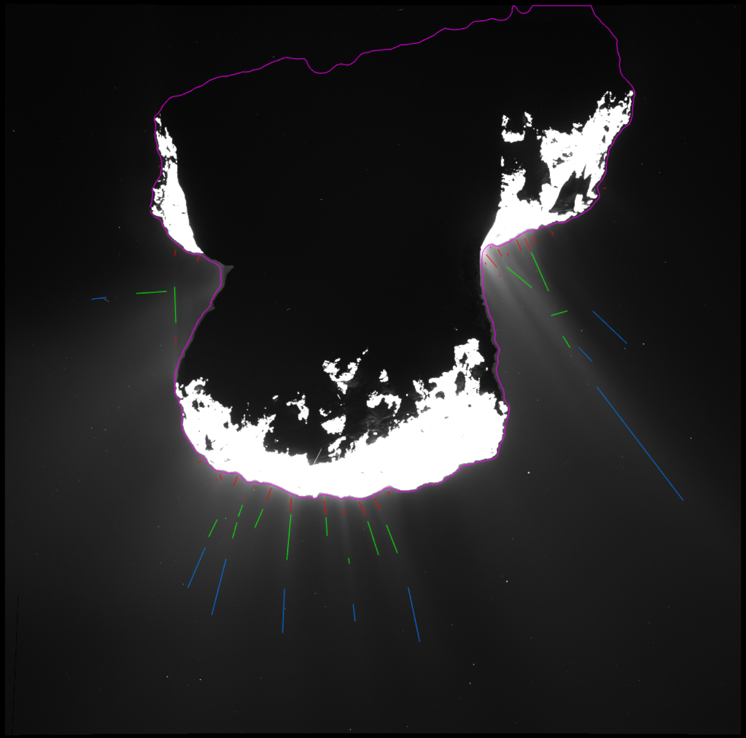

Plume Detection – Plume Vectorization

1. Group nearby pixels together
2. Run Random Sample Consensus (RANSAC) robust line-fitting algorithm on each group, weighted by intensity
3. Apply prior knowledge constraints:
   - Assume plumes are more normal to the nucleus contour
   - Extending plume should intersect nucleus contour

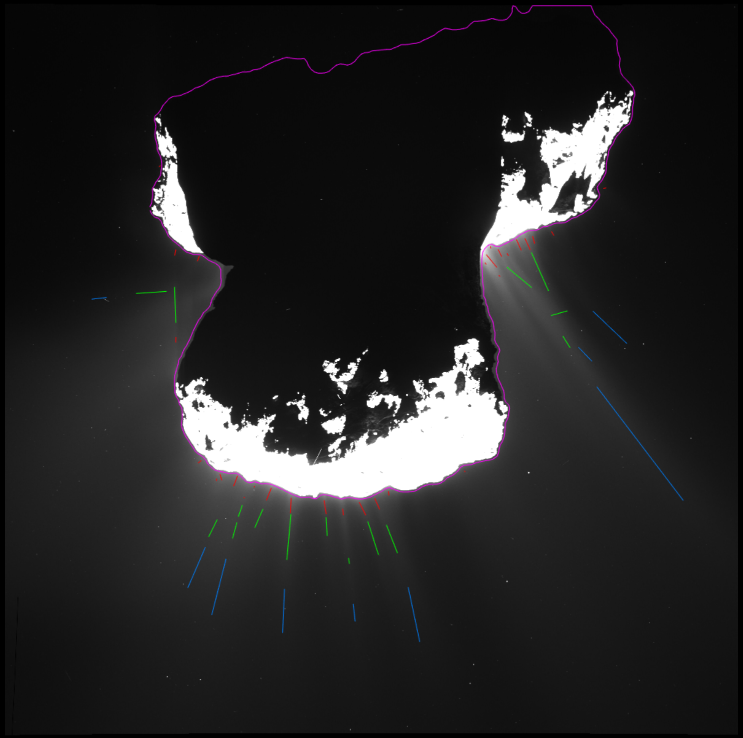

Plume Detection – Plume Extension

- Extend and merge detected plumes until nucleus contour intersection

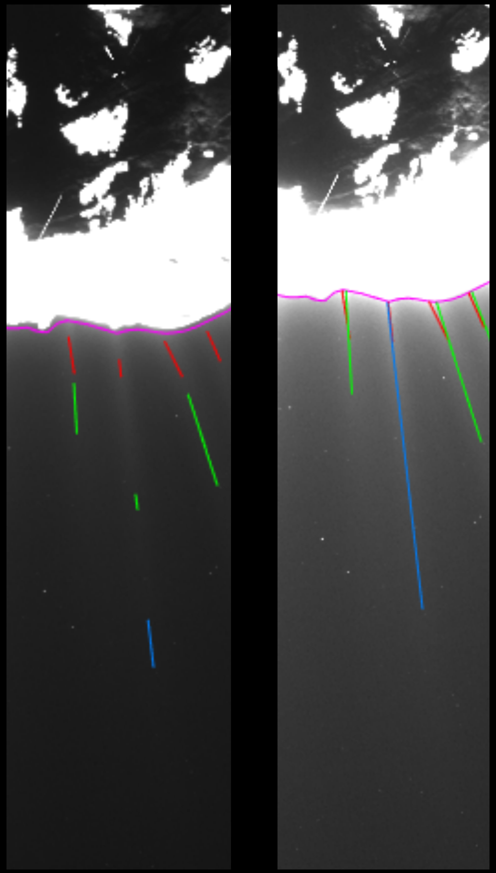

Plume Detection – Plume Extension

- Extend and merge detected plumes until nucleus contour intersection
- Eliminate plumes starting too far from the nucleus contour

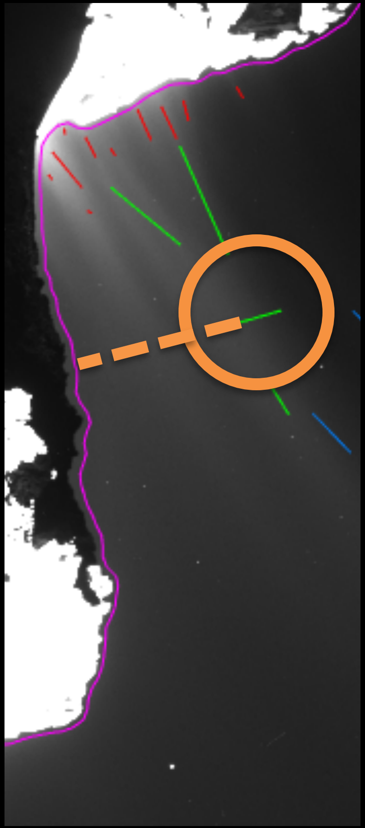

Plume Detection – Plume Extension

- Extend and merge detected plumes until nucleus contour intersection
- Eliminate plumes starting too far from the nucleus contour
- Eliminate plumes nearly tangent to nucleus contour

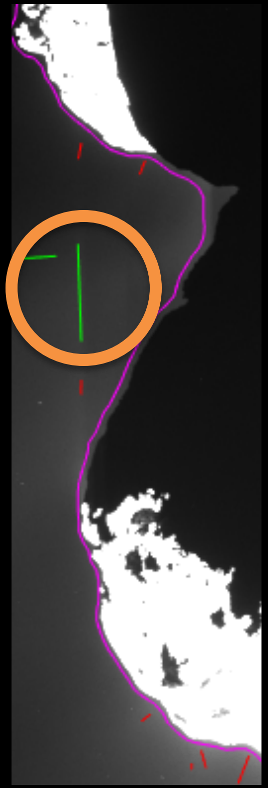

Plume Detection – Plume Extension

- Extend and merge detected plumes until nucleus contour intersection
- Eliminate plumes starting too far from the nucleus contour
- Eliminate plumes nearly tangent to nucleus contour

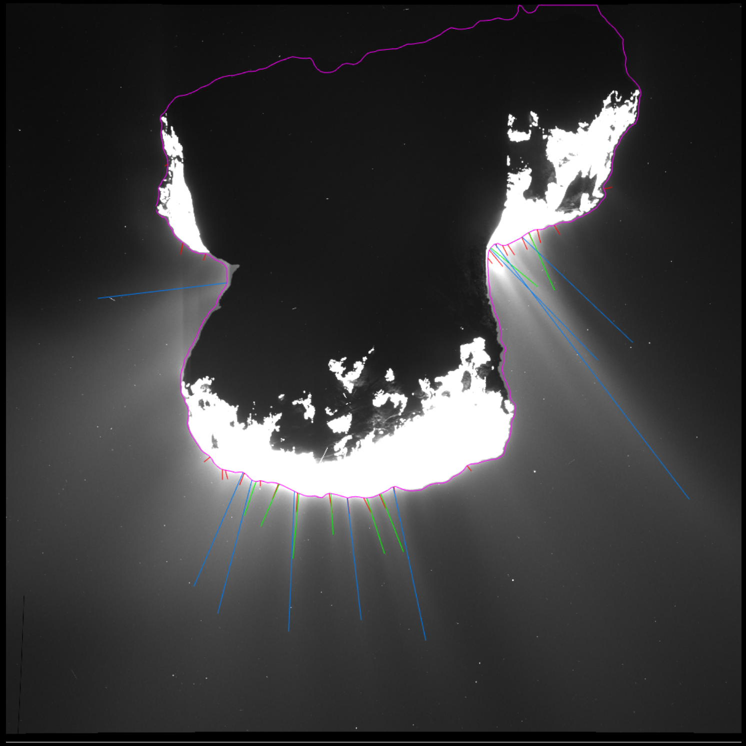

Plume Detection – Example Dataset

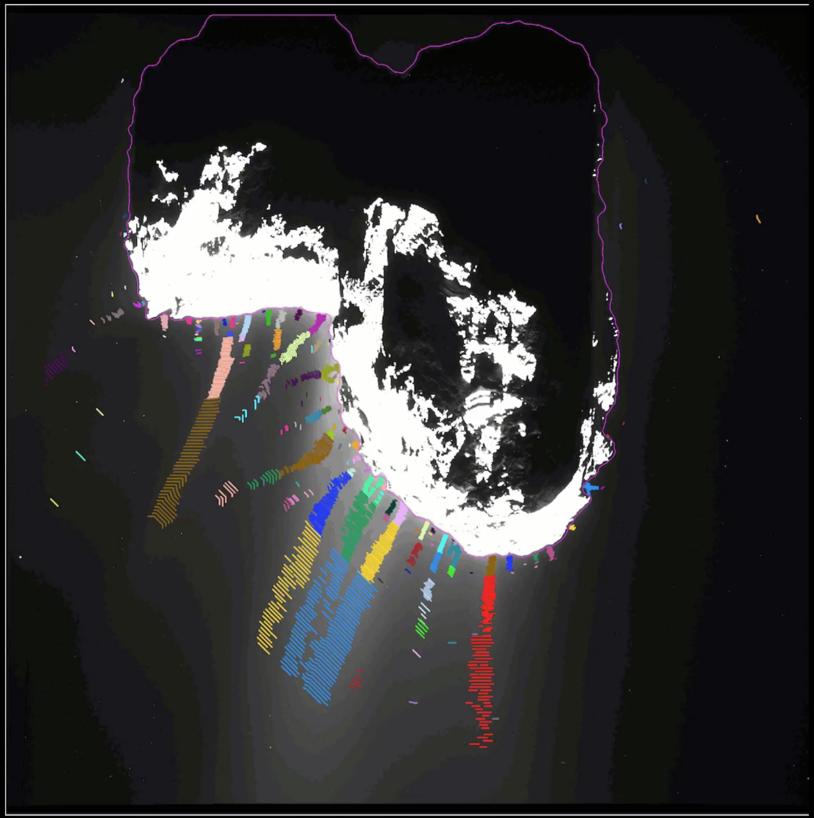

Plume Detection Pipeline

1. Outline the nucleus body
2. Detect candidate plume features in the image
3. Convert into plume vectors
4. Correlate across time

Plume Detection – Cross-Frame Correlation

- Given multiple images taken at different times, can we track the same plume?
- Correlation Pipeline:
    1. Transform plume into C-G frame
    2. Points are mean of a distribution with uncertainty along the observer direction
    3. Calculate the symmetric Kullback-Leibler divergence [2] of all pairs
    4. Produce a ranked list

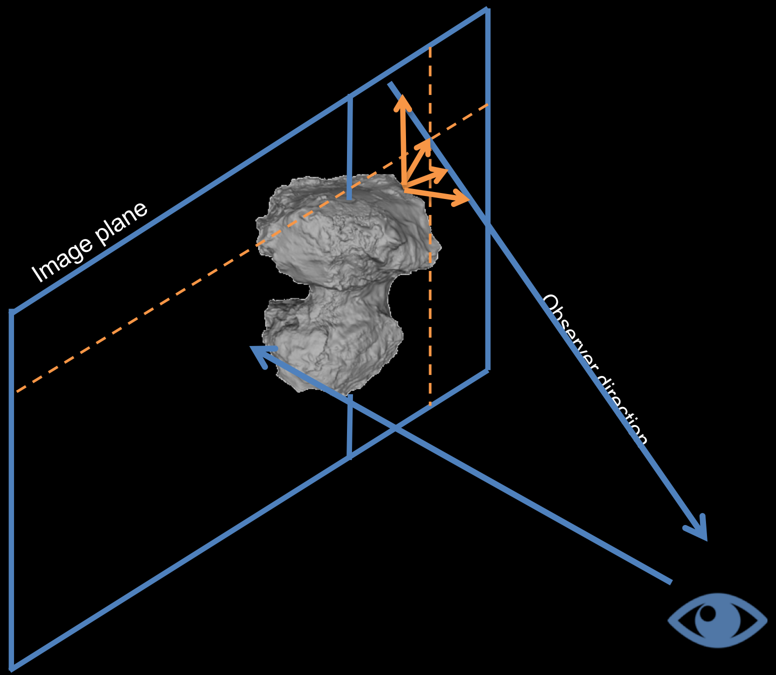

Plume Detection – Cross-Frame Correlation

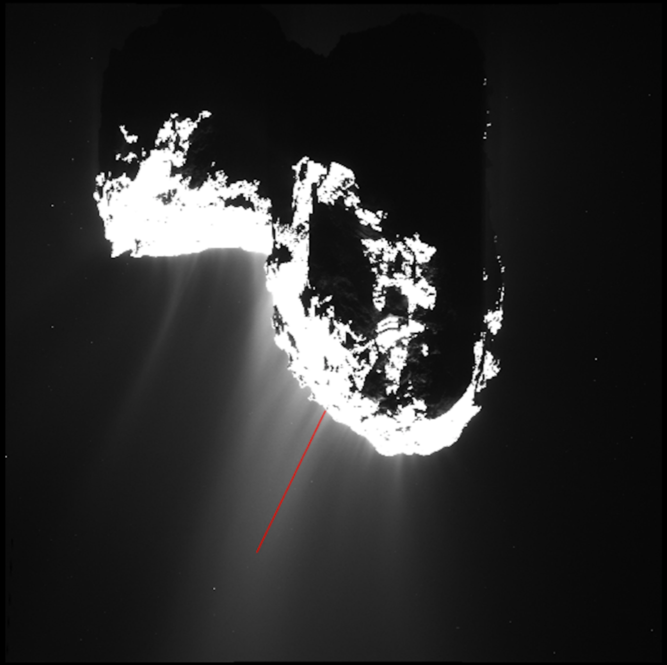

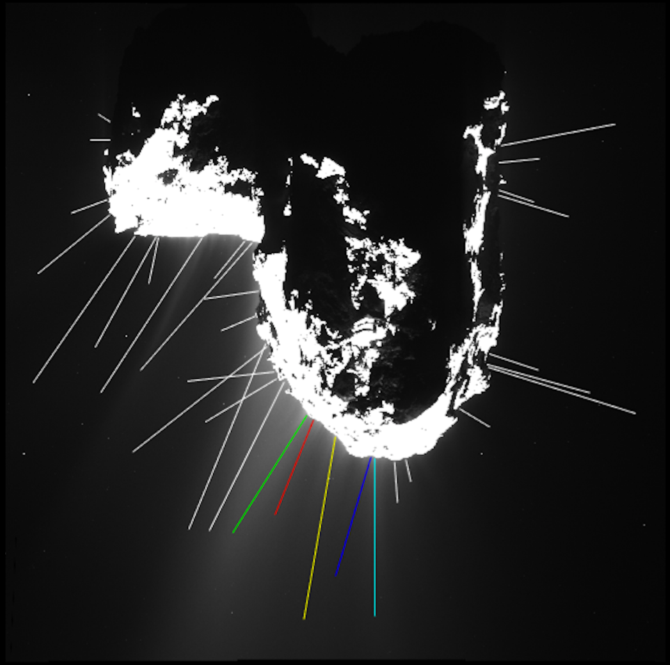

Processing Times – Not Yet Ready for Flight

- 5000 iterations of active contours, ~1 hour per image
- Flexibility in our multi-step approach:
  - Processing time correlated with fidelity of contour
    - Can start with rough shape model
    - More accurate shape model/pointing information -> reduced processing time
    - High fidelity shape model and pointing information could remove the need for image processing
  - Can reduce processing times by focusing on a subset of the image to track a single plume across time

References

1. T. F. Chan, L. A. Vese, Active contours without edges. IEEE Transactions on Image Processing, Volume 10, Issue 2, pp. 266-277, 2001
2. Kullback, S.; Leibler, R.A. (1951). "On information and sufficiency". Annals of Mathematical Statistics. 22 (1): 79–86. doi:10.1214/aoms/1177729694. MR 39968.

Autonomous Response

Onboard Scheduling + Execution

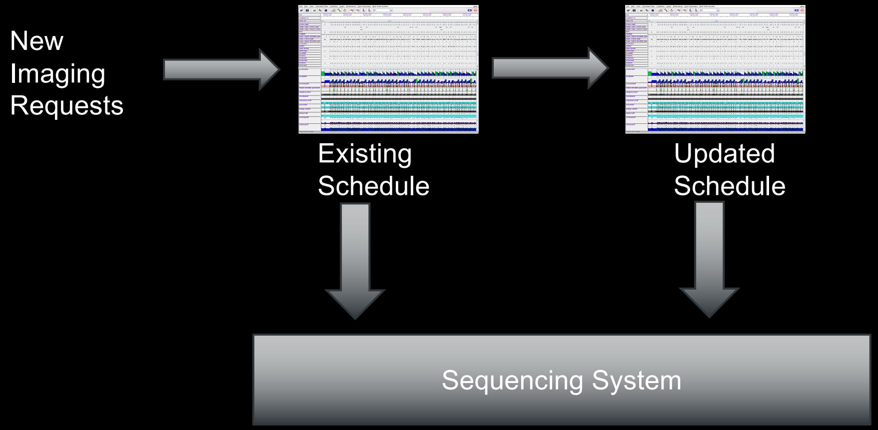

Timeline Management

- Suite of Timeline modeling constructs
  - Finite, infinite states,
    depletable, non-depletable, integral resources
- Typical model elements
  - Power, Energy
  - Data Rates, Volumes
  - Instrument modes
  - Communications systems modes
  - Groundstation views
- Extensive track record for many space mission types
  [Chien et al. 2012, SpaceOps] for further details.

Timeline Management – EO1

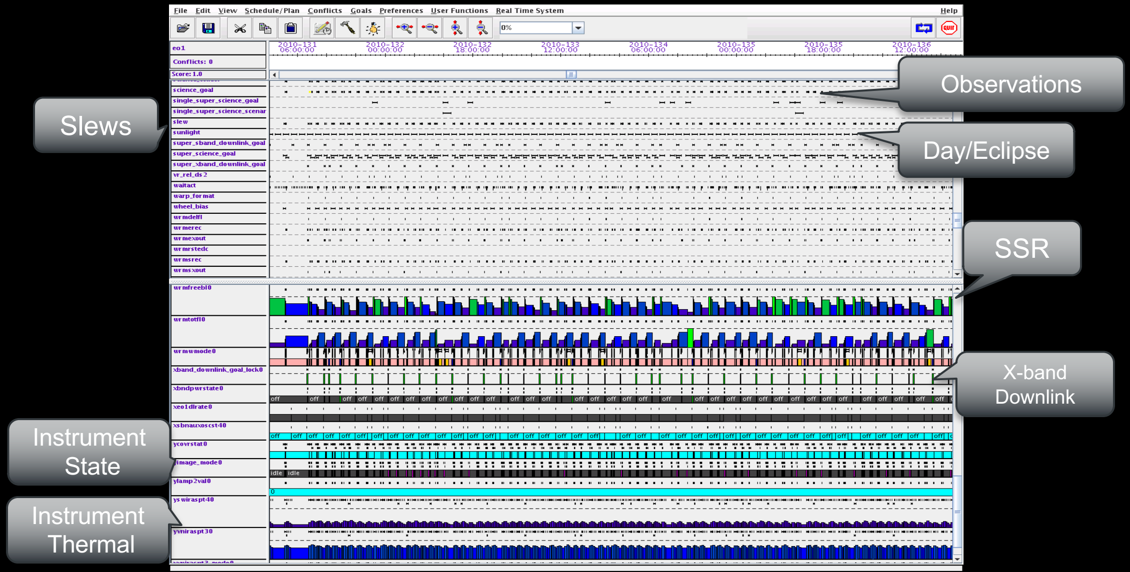

- ASE/EO-1: 1 week observation plan
  See [Chien et al 2005 JACIC, 2010 ICAPS] for further details.

Timeline Management - Rosetta

MTP 006 01 Aug – 01 Sep 2014: 32 days, 2027 observations, 2160 pointings and slews, 63 science campaigns, 10,000's constraints checked and over 1400 downlink dumps. See [Chien et al 2015 IJCAI] for further details.

Scheduling = Iterative constraint satisfaction

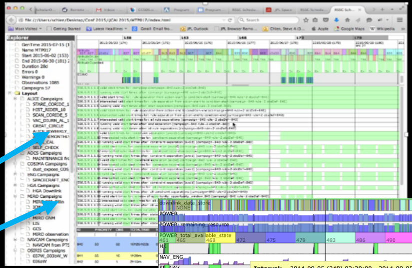

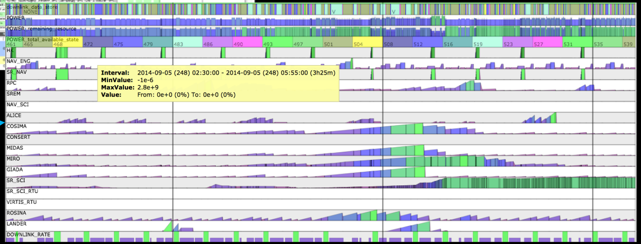

- Iteratively search for temporal placements that satisfy constraints
  - Complicated by interdependency of observations
- Timing Constraints, Geometric Constraints
- Resource Constraints
- Pointing Timeline

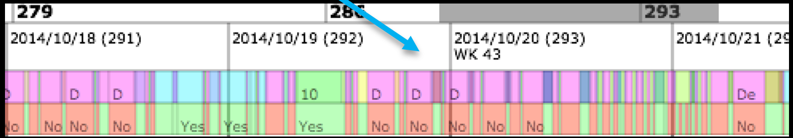

See [Rabideau et al. 2016 ICAPS-Spark for further details)

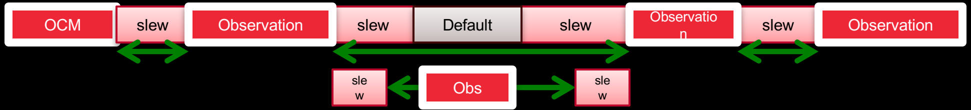

Deep Space: Geometric Computation

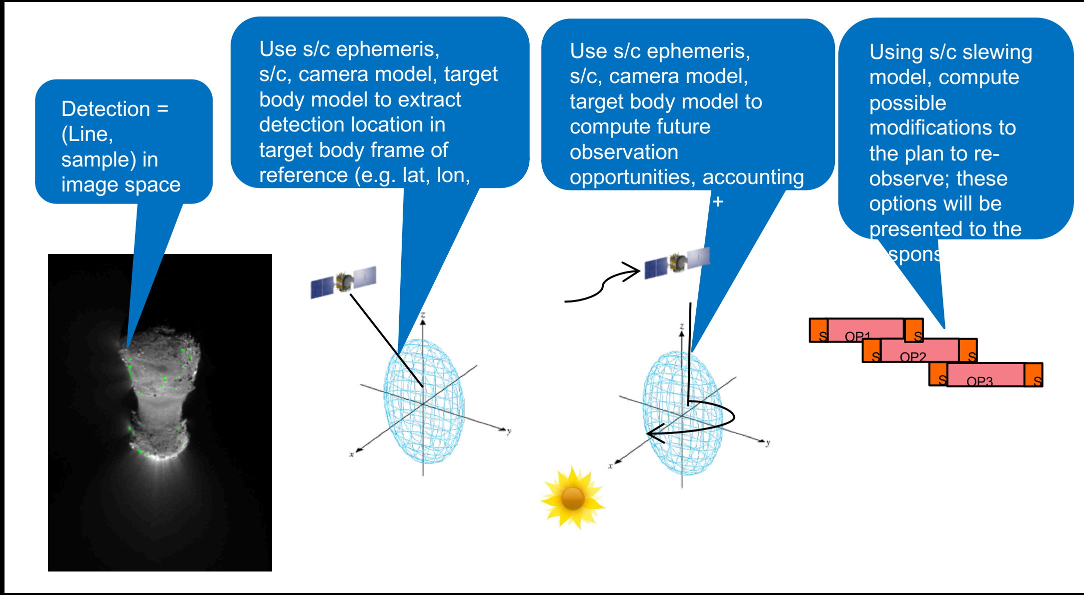

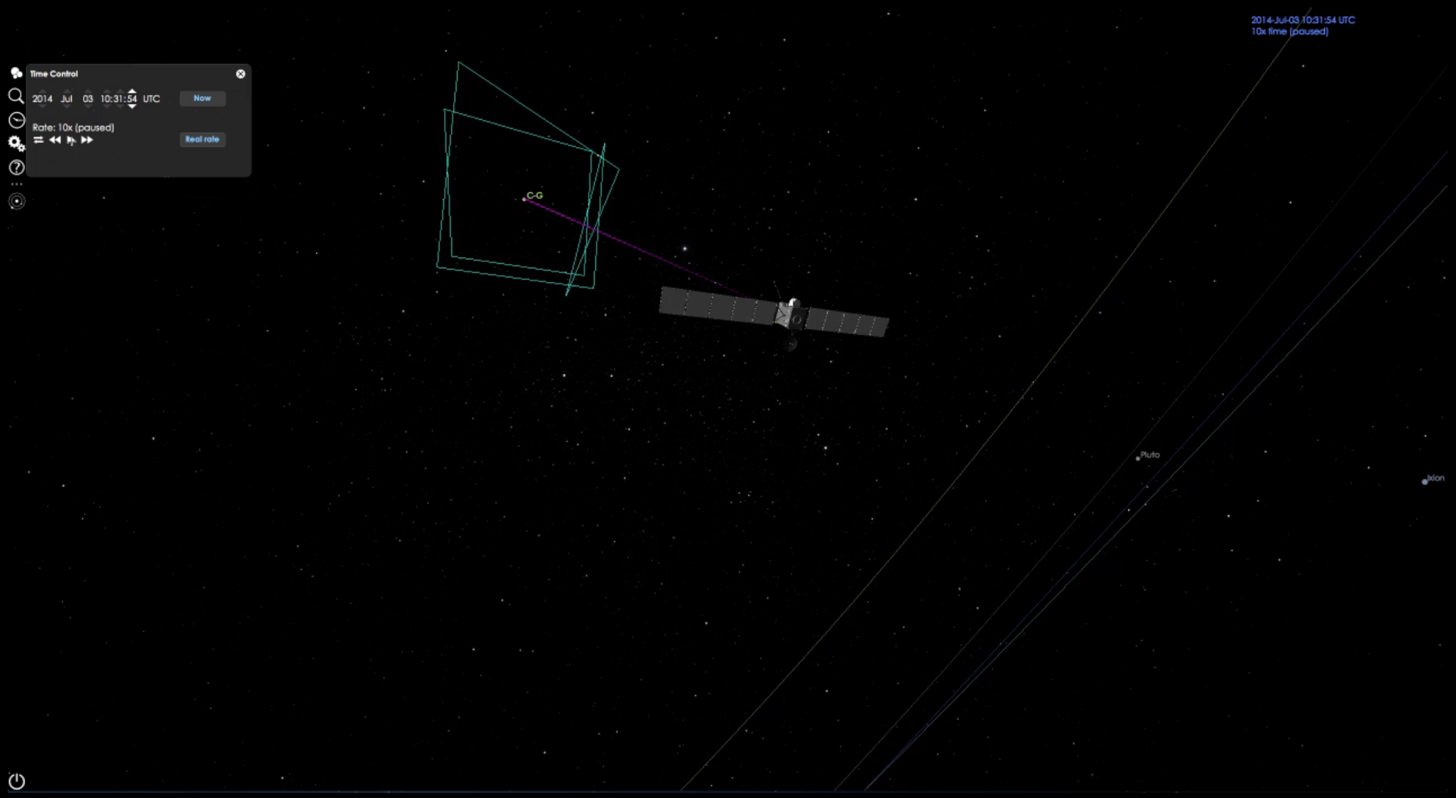
Time Control
2014 Jul 03 10:31:54 UTC
Now
Rate: 10x (paused)
Real rate
C-G
Pluto
Ixion
2014-Jul-03 10:31:54 UTC
10x time (paused)

Conclusions

- Agile Science Technology enables onboard data analysis and response to enhance space missions

- Agile science has already flown on several missions
  - Used in several missions: ASE/EO-1, IPEX, MER, MSL
  - In development for future missions (NEA Scout)

- Several prototypes have been carried to flight software maturity (Agile Science Flyby)

- Future missions can use these technologies to enable new types of science.

Acknowledgements

- Julie Castillo-Rogez and Christophe Sotin for their advice, guidance, and championing the the 4x RTD SI Agile Science

- Julie Castillo-Rogez for her support as NEA Scout Project Scientist

- Art Chmielewski and Claudia Alexander for their support on US-Rosetta

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov